(2)

A Paradigm for Robust Geometric Algorithms*

John E. Hopcroft
Peter J. Kahn

TR 89-1044
October 1989

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*Work on this paper was supported in part by NSF Grant DMC-86-17355, NSF Grant DMS-87-02070, ONR Grant N00014-86-0281, ONR Grant N00014-88-0591, and the US Army Research Office through MSI, Cornell University.

A PARADIGM FOR ROBUST GEOMETRIC ALGORITHMS

JOHN E. HOPCROFT
Department of Computer Science, Cornell University
Ithaca, New York 14853

and

PETER J. KAHN
Department of Mathematics, Cornell University
Ithaca, New York 14853

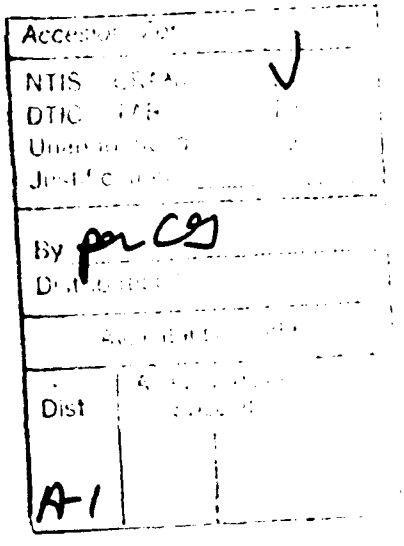

Work on this paper was supported in part by NSF Grant DMC-86-17355, NSF Grant DMS-87-02070, ONR Grant N00014-86-0281, ONR Grant N00014-88-0591, and the US Army Research Office through MSI, Cornell University.

A paradigm for robust geometric algorithms

1. Introduction

Geometrical algorithms in use in computer aided design systems today often fail for numerical reasons. The cause of these failures usually can be traced to logical decisions such as "branch on zero" that depend on the results of numerical calculations. Numerical inaccuracies, introduced either in the initial data or in the finite-precision arithmetic that is used, may result in a set of logical decisions that are inconsistent. This loss of logical consistency usually proves fatal to the algorithm.

Our interest is in geometric algorithms that are robust in the sense that they are provably immune to such potential problems. This paper explores a paradigm that should have wide applicability for producing robust algorithms, and it applies the paradigm to the task of intersecting two convex polyhedral objects. The paradigm was previously used [HHK] in an implementation of an algorithm for intersecting polyhedral objects that was substantially more robust than algorithms implemented earlier. This paper is a first step in developing a mathematical framework that justifies the underlying ideas in this implementation. In particular, an important tool in this work is a method of manipulating embedded polyhedra in ways consistent with their topology.

1.1 Correctness paradigm

The computer representation of a geometrical object consists of two types of data: symbolic and numerical. In the case of a polyhedron, the symbolic data consists of names of vertices, edges and faces along with edge-vertex and face-edge incidence relations. Since symbolic data does not degrade with computation, a program involving only symbolic data should not fail unless there is a flaw in the program itself.

The numerical data usually consist of approximations to the real values of various entities. In the case of a polyhedron, it consists of the coefficients of the face equations

represented to some precision. For a program that performs only numerical computations, an argument based on continuity can be used to establish that the result is an approximation to the correct answer. Of course, such a program may fail for a variety of obvious reasons such as division by zero or an attempt to find the square root of a negative number. These errors can usually be handled in a straightforward manner. A more subtle type of error occurs when there are conversions of numerical data to symbolic data. Typically such a conversion takes place when the result of a numerical computation is used in a branch on zero operation. If the numerical value is very close to zero, one must decide if the quantity is actually zero with the difference due to round-off error. If only one such conversion from numerical data to symbolic data occurs, a slight perturbation of the numerical input data should force the result of the numerical computation to agree with the conversion. Thus, no inconsistency is likely to arise from either decision. However, a program may contain a sequence of such conversions that are logically dependent. Thus, inaccuracies in initial data or in numerical computations may result in inconsistent conversions, and once the symbolic data is inconsistent, some property which is essential for the correctness of the algorithm may fail to hold. Indeed, this is the difficulty that causes many implementations of geometrical algorithms to fail. The solution to this problem is to structure the algorithm, so that all conversions are logically independent. A major contribution of this paper is that an algorithm for an important problem such as the intersection of convex polyhedra can be structured in such a manner.

1.2 A simple example

Consider the following algorithm for intersecting two polygons P_1 and P_2. Intersect each edge of P_1 with each edge of P_2. Use the intersection points to partition edges into segments. Discard each edge segment of P_1 that does not lie in P_2 and each edge segment of P_2 that does not lie in P_1. Assemble the remaining segments into polygons.

If the above algorithm is applied to the two polygons of Figure 1, numerical round-off may result in determining that edge (a,b) intersects edge (d,e) in some point h but that edge

(b,c) fails to intersect edge (d,e). In assembling edges into polygons, it is likely that both segments (d,h) and (h,e) have been discarded as lying outside P_1. In this case there is no second edge at vertex h to associate with (a,h). If by chance (h,e) is classified as inside P_1, then the above difficulty will surely cause problems at vertex e. In either case, the algorithm is likely to return a structure that is not a polygon.

In this simple example, it is clear that the decision as to whether (a,b) intersects (d,e) and the decision as to whether (b,c) intersects (d,e) are logically dependent. The solution to the difficulty is to modify the computation so that all conversions from numerical data to logical data are logically independent. In the above example, this can be done by determining on which side of (d,e) b lies, and then using symbolic reasoning to determine whether (a,b) and (d,e) intersect and whether (b,c) and (d,e) intersect.

The range of problems for which algorithms can be written so that all conversions from numerical data to symbolic data are independent is quite broad and as we shall see includes the intersection of convex polyhedra. It is important to determine the extent of the class of problems for which this can be done. Problems outside the class can be classified by the power of a theorem prover needed to establish that a given set of conversions is or is not independent.

Figure 1: Intersecting two polygons.

1.3 Definition of correctness

A definition of correctness for geometric algorithms must take into account the fact that numerical input data only approximates the real data and thus may be inconsistent with the logical input data. The following definition of correctness, adopted in [HHK], allows for this possible input inconsistency. It depends on suitable notions of approximation, computer

representation, and geometric entity or object, which we discuss in §1.4.

We say that an algorithm for an operation $\circ$ is correct if, for any representations R_1, R_2, $\cdots$ it produces a representation R such that there exists entities M_1, M_2, $\cdots$ and M such that

(1) the numerical data given in R and R_i are approximations to the numerical data given in M and M_i,

(2) the logical data in R and R_i agree with those in M and M_i,

(3) $\circ(M_1, M_2, \cdots) = M$.

Note that correctness of an algorithm in the above sense is not sufficient when the algorithm is used as a subroutine in a more general algorithm. Consider an algorithm for intersecting line segments that reports an intersection only when the line segments themselves or a slight perturbation intersect, and reports a nonintersection only when the line segments themselves or a slight perturbation do not intersect. Such an algorithm is sufficient for a single instance of testing whether two line segments intersect. However, the same algorithm may not be satisfactory in a more general setting. For example, the edge intersection algorithm in the earlier polygon intersection algorithm may be correct by the above definition but fail to work satisfactorily in the polygon intersection setting.

1.4 Computer representations

A geometrical object such as a polyhedron is regarded as a mathematical structure consisting of vertex-edge and edge-face incidences along with equations for the faces. The equations for the faces have coefficients that are real numbers. The geometric incidences determined by these equations match the listed vertex-edge and edge-face incidences. A *computer representation* or simply a *representation* of the object is a data structure presenting the incidence relations and face equations. However, in the computer representation coefficients of the face equations are approximated by finite precision numbers. The vertex-edge and edge-face incidences implied by the given face equations need not agree completely with the

incidence information in the representation. In particular, where the incidence information requires four or more planes to meet at a vertex, the actual geometrical data may imply that the planes meet in a more complex structure consisting of several vertices connected by a tree of short edges (cf. §4.1). The incidence information requires this structure to collapse to a single vertex, thus imposing a condition on the faces that is not stable under small perturbation of the face equations. Representations are discussed further in §4.1.

1.5 Conceptual overview

In this paper the paradigm for robust computation is applied to the problem of intersecting two convex polyhedra. The intersection is carried out by intersecting one of the polyhedra with each half space defining the other. A plane defining a half space may cut the first polyhedron so that certain vertices are on the plane or are so close to the plane that we cannot numerically decide whether they are or are not on the plane. In this case we will declare each such vertex to lie on the plane. We need to prove that we can slightly perturb the face equations of the polyhedron so that the vertices claimed to be on the plane indeed are on the plane. The difficulty here is that faces incident to a high-degree vertex may no longer meet in a point after the perturbation. Thus, we need to develop tools that will allow us to perturb face equations of polyhedra without creating small features. One of these tools is an isomorphism between certain stressed planar graphs and convex polyhedra.

Consider a planar graph whose exterior vertices are fixed to form a convex polygon and whose interior edges behave like springs. The interior vertices will come to rest at an equilibrium position where the forces at each interior vertex sum to zero. Suppose that we can then apply stresses to the edges of the exterior polygon so that the forces at the exterior vertices also sum to zero. The graph is then said to be in equilibrium. There is a correspondence between such graphs in equilibrium and convex polyhedra. It is easier to perturb the vertices of the graph by adjusting the spring constants and thereby perturbing the location of the vertices of the associated polyhedron than it is to perturb the vertices of the polyhedron directly.

In Section 3 we develop the correspondence.

Certain algorithms that deal with polyhedra require that the polyhedron have no small features; e.g., two vertices extremely close together. Such is the case with our algorithm for intersecting a polyhedron with a half space. However, intersecting a polyhedron with a half space may produce a polyhedron with small features. Thus, to iterate the intersection algorithm we need a method to remove small features after each intersection operation. Two options for removing small features are discussed in §4.5.

In Section 5 we apply the isomorphism of Section 3 to the problem of intersecting a convex polyhedron with a half space. For technical reasons, the construction and proof are simpler in the special case where the polyhedron has a triangular face and the plane defining the half space cuts this triangular face into two reasonable size pieces. The general case is briefly described in 5.1.

1.6 Statements of results

There are four major contributions in this paper.

- A paradigm for constructing robust geometrical algorithms.
- A demonstration that decisions based on numerical calculations can be made in a consistent manner for a nontrivial problem: namely the intersection of a convex polyhedron and a half-space.
- An isomorphism between weighted graphs and polyhedra that can be used to manipulate the geometrical embedding of polyhedra in ways consistent with their topology.
- A robust algorithm for intersecting a convex polyhedron and a half space that should never fail for numerical reasons. Under suitable assumptions on precision, the algorithm is correct in the sense of §1.3.

2 Background on convex stressed graphs

In this section we introduce the concept of an equilibrium stress on a planar graph following definitions in [C] and [W]. Given an abstract planar graph with a set of positive stresses, together with a peripheral cycle of the graph embedded as a convex planar polygon, there is a unique extension to a convex embedding of the entire graph satisfying a given equilibrium condition at each interior vertex (cf. Theorem 2.1.3 and Theorem 2.1.4). The importance of this result is that there is a one to one correspondence between sets of positive stresses and planar, convex embeddings of the given graph with fixed periphery.

2.1 Equilibium stresses on planar graphs

Let Γ denote an abstract planar, 3-connected graph. The set of such graphs coincides with the set of 1-skeleta of convex polyhedra [G, p.235]. A mapping of the n vertices of Γ to points in the plane $H = \{(x,y,z) | z = 1\} \subseteq \mathbf{R}^3$ determines a *realization* of Γ. Let us denote the n-tuple of these points $p = (p_1, p_2, \cdots, p_n)$ and the realization $\Gamma(p)$. If the points are distinct and the open edges disjoint, then we say that $\Gamma(p)$ is *embedded*. An embedded graph $\Gamma(p)$ is *convex* if its faces are convex and its periphery is a convex k-gon. Edges not on the periphery are called *interior* edges. It will be useful to select a fixed, abstract, peripheral k-gon in Γ (cf. [T]) and to denote it by Π. In realizations $\Gamma(p)$ of Γ, we may denote the corresponding realization of Π by $\Pi(p)$. We assume, unless stated otherwise, that vertices are numbered so that $p_1, p_2, \cdots, p_k$ correspond to the vertices of $\Pi(p)$.

A *stress* on Γ is a collection $\{\omega_{ij}\}$ of real numbers, $i, j = 1, 2, \cdots, n$, $i \neq j$, satisfying a symmetry and a vanishing condition:

$$\begin{aligned} &(a) \quad \omega_{ij} = \omega_{ji}, && \text{for all } i,j, \quad i \neq j \\ &(b) \quad \omega_{ij} = 0, && \text{unless } p_i \text{ is connected to } p_j \text{ by an edge.} \end{aligned} \tag{1}$$

If each ω_{ij} corresponding to an edge of Γ is strictly greater than zero or less than zero according as the edge is interior or peripheral, we say that $\{\omega_{ij}\}$ is *convex*.

For a given choice of $p = (p_1, p_2, \cdots, p_n)$, the so-called *equilibrium condition* at p_i

$$\sum_{j=1}^{n} \omega_{ij}(p_i - p_j) = 0 \tag{2}$$

may be satisfied for all p_i, in which case we call $\{\omega_{ij}\}$ an *equilibrium stress* and $(\{\omega_{ij}\}, \Gamma(p))$ a *stressed graph*. If (2) is satisfied only when p_i is an "internal" vertex, we call $\{\omega_{ij}\}$ a *restricted equilibrium stress* and $(\{\omega_{ij}\}, \Gamma(p))$ a *restricted stressed graph*. For convenience, we also define a *restricted convex stress* on Γ to be a stress $\{\omega_{ij}\}$ on Γ such that each *internal* ω_{ij} is strictly greater than zero with no condition on the peripheral stresses.

Recall that the vertices p_i belong to H, so that in standard coordinates they are of the form $p_i = (x_i, y_i, 1)$. Sometimes it will be convenient to use instead the coordinates (x_i, y_i). This clearly results in no essential change of information.

Following Connelly [C], we start with a fixed restricted convex stress $\{\omega_{ij}\}$ for Γ, and we define an associated *energy function*.

$$E(p) = \frac{1}{2} \sum_{1 \le i < j \le n} \omega_{ij} |p_i - p_j|^2, \tag{3}$$

where $p = (p_1, p_2, \cdots, p_n)$ ranges over H^n. Clearly, E is a homogeneous quadratic function of the coordinates $x_1, y_1, \cdots, x_n, y_n$.

We now want to hold $x_1, y_1, \cdots, x_k, y_k$, (i. e., $p_1, \cdots, p_k$) fixed. Thus, E becomes a nonhomogeneous quadratic function of the remaining variables: $E(p) = F(x_{k+1}, \cdots, y_n)$. Set $z = (x_{k+1}, \cdots, y_n)$. Then we may write

$$F(z) = Q(z) + L(z) + F(0),$$

where $Q(z)$ (resp., $L(z)$, $F(0)$) is the sum of the quadratic (resp., linear, constant) terms of F.

Define a bilinear form $B(z', z'')$ by the equation

$$B(z', z'') = Q(z' + z'') - Q(z') - Q(z'').$$

Set $m = n - k$ and define the $2m \times 2m$ matrix $\tilde{B}$ by the rule

$$B(z, w) = w^t \cdot \tilde{B} \cdot z.$$

Similarly, define the (column) 2m - tuple $\tilde{L}$ by

$$L(w) = w^t \cdot \tilde{L}.$$

Then the gradient of F is given by

$$\nabla F(z) = \tilde{B} \cdot z + \tilde{L}.$$

From equation (3) above, we can immediately compute that the $(2i - 1)^{st}$ and $2i^{th}$ coordinates of $\nabla F(p_{k+1}, \cdots, p_n)$ are given together by

$$\sum_{j=1}^{n} \omega_{ij}(p_i - p_j), \text{ for } i = k + 1, \cdots, n.$$

Thus, we may immediately conclude the following:

2.1.1 Lemma: The following are equivalent:

(a) $\{\omega_{ij}\}$ is a restricted equilibrium stress for $\Gamma(p)$.

(b) $z = (p_{k+1}, \cdots, p_n)$ is a critical point of F (i. e., $\nabla F(p_{k+1}, \cdots, p_n) = 0$).

(c) $\tilde{B} \cdot z = -\tilde{L}$. □

We note here that after a suitable permutation of rows and columns, the $2m \times 2m$ matrix $\tilde{B}$ actually has the simpler form

$$\begin{bmatrix} C & O \\ O & C \end{bmatrix}$$

for a certain $m \times m$ matrix C (see Lemma 2.2.1(b)).

2.1.2 Lemma: F has a unique critical point, at which it achieves its minimum.

Proof: The proof is modeled on a proof in Connelly [C]. Suppose that $z = (p_{k+1}, \cdots, p_n)$ is large. Then p_j is large for some $j > k$. This implies that in the shortest edge-path connecting p_j to some $p_i, i \leq k$, some edge is large. But the coefficient in (3) corresponding to this edge is strictly greater than zero, and so $E(p) = F(z)$ is large Thus, for sufficiently large $\alpha > 0, |z| > \alpha$ implies that $F(z) > F(0)$. It follows that F must achieve its minimum on

$\{z \mid |z| \leq \alpha\}$.

To prove uniqueness, suppose F achieves its minimum at z' and that z'' is a critical point of F. For real t, define $f(t) = F(tz' + (1 - t)z'')$. This function is at most quadratic in t. Hence, its derivative $f'(t)$ is at most linear. By assumption, $f'(t)$ satisfies $f'(0) = f'(1) = 0$. Thus, f is constant. But then $tz' + (1 - t)z''$ must be constant since $F(z)$ is large for large z. Therefore $z'' = z'$. □

2.1.3 Theorem: Choose a convex k-gon in H with consecutive vertices $q_1, q_2, \cdots, q_k$. Suppose $\{\omega_{ij}\}$ is any restricted convex stress on Γ. Then there exists a unique realization $\Gamma(p)$ of Γ satisfying:

(1) $p_i = q_i, i = 1, 2, \cdots, k$,

(2) $\{\omega_{ij}\}$ is a restricted equilibrium stress for $\Gamma(p)$, and

(3) the points $p_{k+1}, p_{k+2}, \cdots, p_n$ depend smoothly on the ω_{ij}, $i \geq k+1$ or $j \geq k+1$.

Proof: Define the energy function E as in (3). Set $p_i = q_i, i = 1, 2, \cdots, k$, and define F as above. By Lemma 2.1.2, F has a unique critical point $z = (p_{k+1}, \cdots, p_n)$, which implies, by Lemma 2.1.1, that $\{\omega_{ij}\}$ is a restricted equilibrium stress for $\Gamma(p)$, as desired.

Uniqueness of the realization follows from the uniqueness portion of Lemma 2.1.2.

Finally, to see that $z = (p_{k+1}, \cdots, p_n)$ depends smoothly on $\{\omega_{ij} \mid i \geq k+1$ or $j \geq k+1\}$, we consult equation (c) of Lemma 2.1.1: $\tilde{B}z = -\tilde{L}$. As we shall see in Lemma 2.2.1 below, $\tilde{B}$ and $\tilde{L}$ depend smoothly on the above stress constants. To complete the proof, then, it remains to show that the matrix $\tilde{B}$ is invertible, for then,

$$z = -(\tilde{B}^{-1})\cdot\tilde{L},$$

and smooth dependence of z follows.

But the invertibility of $\tilde{B}$ follows immediately from the equivalence 2.1.1(b) $\Leftrightarrow$ 2.1.1(c), together with the uniqueness part of Lemma 2.1.2. □

Tutte [T] obtains a special case of Theorem 2.1.3, making use of an inductive computation in an earlier paper [BSST]. The proof of the general case presented here is based on some ideas of Connelly [C] and leads nicely to subsequent calculations.

2.1.4 Theorem (Tutte [T]): Suppose that $\Gamma(p)$ is a realization of Γ such that $\Pi(p)$ is a convex k-gon. Suppose that there exists a restricted convex stress on Γ which is a restricted equilibrium stress for $\Gamma(p)$. Then $\Gamma(p)$ is convex. □

In particular, *the existence of a convex, restricted equilibrium stress for* $\Gamma(p)$ *(with* Γ 3-connected) *implies that* $\Gamma(p)$ *is embedded,* with all nonperipheral vertices $p_{k+1}, \cdots, p_n$ actually interior to $\Pi(p)$. Tutte [T] proves this result in the case that all the interior ω_{ij} equal one, but the proof extends virtually verbatim to our more general case.

Theorem 2.1.3 tells us that if we pin the peripheral vertices at $p_1, p_2, \cdots, p_k$ so as to form a convex k-gon in H, then any assignment of positive stress constants to internal edges of Γ will uniquely and smoothly produce internal vertices, at each of which the equilibrium conditions are satisfied. Theorem 2.1.4 asserts that the resulting $\Gamma(p)$ is convex (hence embedded). Unfortunately, Theorem 2.1.3 (2) above gives us only a restricted equilibrium stress, whereas our later applications require, in addition, equilibrium at each peripheral vertex $p_1, p_2, \cdots, p_k$. Theorem 2.3.1 below shows that we can obtain these additional equilibrium conditions when $k = 3$.

2.2 Some computations

Given a stress $\{\omega_{ij}\}$ on Γ, we follow [C] and define an associated *stress matrix* Ω as follows:

$$\Omega_{ij} = \begin{cases} -\omega_{ij}, & i \neq j \\ \sum_{k=1}^{n} \omega_{ik}, & i = j \end{cases} \qquad i,j = 1, 2, \cdots, n . \tag{4}$$

The equilibrium equations

$$\sum_{j=1}^{n} \omega_{ij}(p_i - p_j) = 0, \quad i = 1,2, \cdots, n \tag{5}$$

may now be rewritten as the pair of matrix equations

$$\begin{aligned} \Omega \cdot X &= 0 \\ \Omega \cdot Y &= 0 \end{aligned} \tag{6}$$

where X (resp., Y) is the column n-tuple $(x_1, x_2, \cdots, x_n)^t$ (resp., $(y_1, y_2, \cdots, y_n)^t$) consisting of the x-coordinates (resp., y-coordinates) of the $p_i = (x_i, y_i, 1)$.

Recall that vertices have been numbered so that the vertices of the peripheral k-gon $\Pi(p)$ are $p_1, p_2, \cdots, p_k$. Then, we write Ω in block form as follows

$$\Omega = \begin{bmatrix} A & B \\ B^t & C \end{bmatrix} \tag{7}$$

where A is $k \times k$ and C is $m \times m$, $m = n - k$. It is easy to see that the equilibrium equations (5) hold for $i = k+1, k+2, \cdots, n$, if and only if

$$\begin{aligned} B^t X' + C X'' &= 0 \\ B^t Y' + C Y'' &= 0 \end{aligned} \tag{8}$$

where $X' = (x_1, \cdots, x_k)^t$, $X'' = (x_{k+1}, \cdots, x_n)^t$ and similarly for Y' and Y''. Thus, (8) expresses the fact that $\{\omega_{ij}\}$ is a *restricted* equilibrium stress on $\Gamma(p)$.

Comparing (8) with Lemma 2.1.1, we can easily conclude the following:

Lemma 2.2.1: (a) Up to a permutation of coordinates,

$$\tilde{L} = \begin{bmatrix} B^t X' \\ B^t Y' \end{bmatrix}$$

(b) Up to a permutation of rows and columns

$$\tilde{B} = \begin{bmatrix} C & O \\ O & C \end{bmatrix}. \quad \square$$

Corollary 2.2.2: Assume that $p_1, \cdots, p_k$ (hence X', Y') are fixed, and that $\{\omega_{ij}\}$ is a restricted convex equilibrium stress on Γ. Then C is invertible, and the vertices $p_{k+1}, \cdots, p_n$ produced by Theorem 2.1.3 (i. e., X'', Y'') may be computed by

$$\begin{aligned} X' &= -C^{-1}B^tX' \\ Y' &= -C^{-1}B^tY' \end{aligned} \tag{9}$$

Proof: The proof of Theorem 2.1.3 shows that $\bar{B}$ is invertible. Hence C is, by Lemma 2.2.1(b) above. Then (9) follows immediately from (8). □

2.3 Finding stress constants for graphs with triangular periphery

2.3.1 Theorem: Suppose that $\Gamma(p)$ is a realization of Γ with $\Pi(p)$ a nondegenerate triangle (i. e., $k=3$). Suppose also that $\{\omega_{ij}\}$ is a restricted equilibrium stress for $\Gamma(p)$. Then, there exists a unique equilibrium stress $\{\omega_{ij}'\}$ for $\Gamma(p)$ satisfying:

(1) $\omega_{ij}' = \omega_{ij}$, if i or $j > 3$.

(2) $\{\omega_{ij}' \mid 1 \leq i, j \leq 3\}$ depends smoothly on $\{\omega_{ij} \mid i \text{ or } j > 3\}$.

(3) If $\{\omega_{ij}\}$ is restricted convex, then $\{\omega_{ij}'\}$ is convex.

Proof: We use the notation and computations in 2.2. Thus, $\{\omega_{ij}\}$ corresponds to a stress matrix

$$\Omega = \begin{bmatrix} A & B \\ B^t & C \end{bmatrix}$$

as in (7), and the x- and y-coordinates of p yield column n-tuples

$$X = \begin{bmatrix} X' \\ X'' \end{bmatrix} \quad \text{and} \quad Y = \begin{bmatrix} Y' \\ Y'' \end{bmatrix}$$

where X' and Y' are column triples, etc. Let $\bar{1}$ denote the column n-tuple $(1,1,\ldots,1)^t$, with $\bar{1}'$ and $\bar{1}''$ the corresponding triple and $(n-3)$-tuple.

Let Ω' denote the stress matrix corresponding to the sought after equilibrium stress $\{\omega_{ij}'\}$. Condition (1) of the theorem requires that Ω' has the form

$$\Omega' = \begin{bmatrix} A' & B \\ B^t & C \end{bmatrix}$$

where A' is a 3×3 matrix to be determined. The equilibrium condition is equivalent to the

equations

$$\Omega' \cdot X = 0$$

$$\Omega' \cdot Y = 0$$

whereas the fact that Ω' has row-sums equal to zero is equivalent to

$$\Omega' \cdot \bar{1} = 0.$$

These yield the following equation involving A':

$$A' \cdot [X', Y', \bar{1}'] + B \cdot [X'', Y'', \bar{1}''] = 0.$$

Since the triangular periphery of $\Gamma(p)$ is nondegenerate, the 3×3 matrix $[X', Y', \bar{1}']$ is nonsingular. Thus, the last equation determines A' uniquely:

$$A' = -B \cdot [X'', Y'', \bar{1}''] \cdot [X', Y', \bar{1}']^{-1}.$$

This demonstrates the existence and uniqueness of an equilibrium stress $\{\omega_{ij}'\}$ on $\Gamma(p)$ satisfying (1) and (2). It remains to demonstrate (3), which we do by a geometric argument.

Note that if equilibrium is satisfied at a vertex, there cannot exist a line through that vertex separating the positively stressed edges from the negatively stressed ones. It follows immediately that, under conditions of equilibrium, if the internal edges meeting a given peripheral vertex are all positively stressed, then the peripheral edges meeting that vertex must be negatively stressed. □

3. Polyhedra and stressed planar graphs

In this section we describe a basic correspondence between certain planar graphs and certain polyhedral surfaces. This correspondence allows us, in some cases of interest, to reduce problems of manipulating polyhedra to problems of manipulating graphs. Our view is that this gives us a modeling tool with potentially wide applicability. We back up this claim in Section 4 with applications to removing small features and intersecting polyhedra with half spaces.

The families of polyhedral surfaces and of stressed planar graphs that we consider have natural structures as vector spaces, with one mapped to the other by a linear isomorphism. The mapping was first noted by Maxwell [M] and the bijectivity property by Crapo and Whiteley [W]. Our description of the mapping differs somewhat from earlier descriptions and has several advantages. It is given by an explicit formula, it is visibly linear, and it uses only well-known, elementary linear algebra. The proof of bijectivity, however, requires a homological argument. This follows easily from our formulation and is given in Appendix B. In fact, this argument is a special case of a quite general homological construction that will appear elsewhere [K].

The authors wish to thank Robert Connelly for a number of helpful conversations and suggestions concerning the material in this section.

3.1 Stressed graphs

Let $\Gamma(p)$ be a 3-connected, planar graph embedded in the plane $H = \{(x,y,z) \mid z = 1\}$, with boundary k-gon $\Pi(p)$. Let $S(\Gamma, p)$ denote the set of equilibrium stresses $\{\omega_{ij}\}$ on $\Gamma(p)$. We may define addition and scalar multiplication of equilibrium stresses by the rules

$$\{\omega_{ij}\} + \{\bar{\omega}_{ij}\} = \{\omega_{ij} + \bar{\omega}_{ij}\}$$

$$c\,\{\omega_{ij}\} = \{c\,\omega_{ij}\}.$$

With these operations, $S(\Gamma, p)$ becomes a vector space.

3.2 Polyhedral surfaces

Let $f_1, f_2, \cdots$ denote the interior faces of $\Gamma(p)$, and let f_0 be the closed region bounded by $\Pi(p)$, which we also call a face. We now describe polyhedral surfaces that "sit over" the graph $\Gamma(p)$. Intuitively, such a surface may be constructed from planar polygonal faces in $\mathbf{R}^3$, one for each face in $\Gamma(p)$. Each face in $\mathbf{R}^3$ projects orthogonally down to its counterpart in $\Gamma(p)$ and matches up on edges with neighboring faces. In order to deal with these surfaces algebraically, we recognize that if f_r is a face of $\Gamma(p)$, then any polygonal face in $\mathbf{R}^3$ that pro-

jects to f_r may be envisioned as the graph of an affine function defined on f_r. Thus, a polyhedral surface may be represented as a certain family of affine functions. We now make this precise.

Recall that $\Gamma(p)$ sits in the plane H. Let $A(H)$ denote the set of all affine functions $H \to \mathbf{R}$. A typical such function a is given by a formula $a(x,y,1) = cx + dy + e$, for certain constants c,d,e. Addition and scalar multiplication of functions make $A(H)$ into a vector space of dimension three. Let $\{a_r\}$ be a sequence of affine functions in $A(H)$, one for each face f_r of $\Gamma(p)$ (including f_0). We say that the sequence $\{a_r\}$ is *piecewise-affine* (P.A.) if the a_r satisfy a suitable compatibility property. Namely, for each edge e_{ij} and incident faces f_r, f_s, we must have

$$a_r \mid e_{ij} = a_s \mid e_{ij} \tag{10}$$

where $a \mid e$ is the restriction of a to e. If we imagine the graphs of $a_r \mid f_r$ and $a_s \mid f_s$ as sitting in $\mathbf{R}^3$ "over" the plane H, then (10) simply announces that these graphs match up over the edge e_{ij}. Thus, the sequence $\{a_r\}$ is a representation of a surface sitting over $\Gamma(p)$.

Figure 2: Polyhedron sitting over plane and the graph of its projection.

We now collect all the piecewise-affine sequences $\{a_r\}$ into a set, denoted $PA(\Gamma, p)$, and note that the usual operations of termwise addition and scalar multiplication of sequences make $PA(\Gamma, p)$ into a vector space. By our remarks above, this vector space may be identified with the space of polyhedral surfaces over $\Gamma(p)$.

Now suppose that $\{b_r\}$ is a P.A. sequence and $b_0 = 0$. This means that in the corresponding polyhedral surface, the face represented by the graph of $b_0 \mid f_0$ actually coincides with f_0. In effect, the polyhedral surface "sits on" f_0 (in H). If $\{a_r\}$ is any P.A. sequence, we can simply subtract a_0 from each a_r to obtain a P.A. sequence sitting on f_0. This translates into the

following simple assertion about the vector space $PA(\Gamma,p)$ of all P.A. sequences. Let $PA(\Gamma,p)_0$ consist of all P.A. sequences $\{a_r\}$ sitting on f_0 (i.e., $a_0=0$). This is a vector subspace of $PA(\Gamma,p)$, as is the set of all constant sequences (i.e., sequences $\{a_r\}$ with $a_r=a_0$, for all r). Together, these subspaces give a direct sum decomposition of $PA(\Gamma,p)$:

$$PA(\Gamma,p) = PA(\Gamma,p)_0 + \{\text{constant sequences}\}.$$

Since a constant sequence is specified by any one term, it may be identified with that term, which is an affine function (an element of $A(H)$). Thus, {constant sequences} may be identified with the vector space $A(H)$, and we have the direct sum decomposition:

$$PA(\Gamma,p)=PA(\Gamma,p)_0+A(H).$$

One caution should be given here. What we have called a polyhedral surface may have self-intersections. It may even flatten out into a plane, as in the case of a constant sequence. Under some fairly obvious conditions, however, a P.A. sequence will give a polyhedral surface which is a topological image of the standard 2-sphere. For such a "nonsingular" surface, we may refer to its interior, and we call the closure of this interior the corresponding *polyhedron*. In our applications, we shall be dealing with convex surfaces, and these are always nonsingular.

3.3 The map Δ: $PA(\Gamma,p)\to S(\Gamma,p)$

We now define a map Δ from piecewise affine sequences on $\Gamma(p)$ to equilibrium stresses on $\Gamma(p)$. The map Δ: $PA(\Gamma,p)\to S(\Gamma,p)$ will be seen to be a linear transformation. Recall that Γ and $p_1,\cdots,p_n$ are fixed. Thus, it suffices for each e_{ij} to express ω_{ij} in terms of the a_r and the vertices $p_1,\cdots,p_n$. Choose an arbitrary reference point $p_*\in H=\{z=1\}$ that is not collinear with any edge of $\Gamma(p)$. Then

$$\omega_{ij}=\varepsilon(r,s)(a_s(p_*)-a_r(p_*))/[p_i,p_j,p_*]\,, \tag{11}$$

where $[p_i,p_j,p_*]$ is the usual triple product in $\mathbf{R}^3$. If p_i,p_j, and p_* are regarded as column

vectors, then $[p_i, p_j, p_*] = \det(p_i, p_j, p_*)$.

The factor $\varepsilon(r,s)$ equals $+1$ or -1 according to certain orientation conventions. Assume $\mathbf{R}^3$ is oriented by the "right-hand rule," so that when viewed from above, a positive rotation in H is counterclockwise. This counterclockwise orientation in H imparts an orientation to each interior face f_r of $\Gamma(p)$, hence an orientation to each edge in the boundary of f_r. We give f_0 a clockwise orientation.

Now consider the given edge e_{ij}, $i<j$. It is incident to precisely two faces, say f_r and f_s. This defines the r and s in (11). We set $\varepsilon(r,s)$ equal to $+1$ if, in the orientation imparted by f_r, e_{ij} points from p_i to p_j. Otherwise, we set $\varepsilon(r,s)$ equal to -1.

Note that the value of (11) is not affected by an exchange of r and s, since this changes the sign of both factors in the numerator. Thus, the right-hand side of (11) depends only on i and j, so that ω_{ij} is well-defined.

This concludes our description of (11).

Now (11) defines ω_{ij} when p_i is adjacent to p_j, and $i<j$. By using equations (1) $(a),(b)$ in §1, we extend this to a definition of ω_{ij} for $1 \le i, j \le n$, $i \ne j$, obtaining a stress $\{\omega_{ij}\}$ on Γ.

3.3.1 Theorem: $\{\omega_{ij}\}$ is an equilibrium stress on $\Gamma(p)$.

We prove this in Appendix A.

Formula (11) shows clearly that $\{\omega_{ij}\} = \Delta\{a_r\}$ depends linearly on $\{a_r\}$, i.e., that Δ is a linear transformation $PA(\Gamma,p) \to S(\Gamma,p)$. Its nullspace is easy to describe.

3.3.2 Theorem: The nullspace of Δ consists of the constant sequences in $PA(\Gamma,p)$; i.e., using our comments in 3.2, nullspace $\Delta = A(H)$.

Proof: Formula (11) implies that $\{a_r\}$ is the constant sequence if and only if the correspond-

ing ω_{ij}'s are all zero. □

Now using the direct sum decomposition in 3.2,

$$PA(\Gamma,p)=PA(\Gamma,p)_0+A(H),$$

we get the following.

3.3.3 Corollary: $\Delta \mid PA(\Gamma,p)_0 \to S(\Gamma,p)$ **is** $1-1$. □

Finally, we state the key result of this section:

3.3.4 Theorem: Δ: $PA(\Gamma,p) \to S(\Gamma,p)$ (or $\Delta \mid PA(\Gamma,p)_0 : PA(\Gamma,p)_0 \to S(\Gamma,p)$) **is onto.**

The proof of this theorem uses some elementary homology theory; it is presented in Appendix A.

Combining Corollary 3.3.3 and Theorem 3.3.4, we have:

3.3.5 Corollary: $\Delta \mid PA(\Gamma,p)_0$ **is an isomorphism** $PA(\Gamma,p)_0 \to S(\Gamma,p)$.

Corollary 3.3.5 gives the $1-1$ correspondence between polyhedral surfaces and stressed graphs announced in our introductory paragraph. Thus, schematically, we have

$$\left\{\begin{matrix}Surfaces\ S\ over\ \Gamma(p)\\ sitting\ on\ f_0\end{matrix}\right\} \to \underbrace{\left\{\begin{matrix}P.A.Sequences\\ \{a_r\}\ s.t.\ a_0=0\end{matrix}\right\}}_{PA(\Gamma,p)_0} \overset{\Delta_0}{\to} \underbrace{\left\{\begin{matrix}Equilibrium\ stresses\\ \{\omega_{ij}\}\ on\ \Gamma(p)\end{matrix}\right\}}_{S(\Gamma,p)} \tag{12}$$

where $\Delta_0 = \Delta \mid PA(\Gamma,p)_0$.

3.4 Varying $\Gamma(p)$

For our purposes we will need to vary the location of $p_1, p_2, \cdots, p_n$ (the "geometry" of $\Gamma(p)$), while fixing Γ (the "topology" of $\Gamma(p)$). We allow $p_1, p_2, \cdots, p_n$ to vary in H, subject to the following restrictions:.

$$(a)\ p_1, p_2, \cdots, p_n \text{ are distinct;}$$

$$(b)\ \Gamma(p) \text{ is a planar embedding of } \Gamma. \tag{13}$$

We modify the spaces in diagram (12) to make explicit the dependence on Γ and p:

$$\left\{\begin{array}{c}\textit{Pairs } (S,\ \Gamma(p)) \\ \textit{where } S \textit{ is a} \\ \textit{surface over } \Gamma(p) \\ \textit{sitting on } f_0\end{array}\right\} \rightarrow \left\{\begin{array}{c}\textit{Pairs}(\{a_r\},\ \Gamma(p)) \\ \textit{s.t.} \{a_r\} \textit{ is a P. A.} \\ \textit{sequence relative to} \\ \Gamma(p),\ \textit{with } a_0 = 0\end{array}\right\} \overset{\Delta_0}{\rightarrow} \left\{\begin{array}{c}\textit{Pairs}(\{\omega_{ij}\},\ \Gamma(p)) \\ \textit{where } \{\omega_{ij}\} \textit{ is an} \\ \textit{equilibrium stress} \\ \textit{on } \Gamma(p)\end{array}\right\}$$

Each of these sets has a natural topology. For example, suppose Γ has e edges. Then, each pair $(\{\omega_{ij}\}, \Gamma(p))$ corresponds to an $(e+3n)$-tuple $(\ \cdots, \omega_{ij}, \cdots, p_1, \cdots, p_n)$, so that $S(\Gamma, p)$ may be considered a subset of R^{e+3n}. Similarly for the other sets. The left-hand set and $PA(\Gamma, p)_0$ are essentially two versions of the same thing, as described in §3.2. The left-hand arrow simply denotes the transition from one version to another and is thus clearly a homeomorphism. The right-hand arrow Δ_0 is also a homeomorphism, but this requires some explanation.

We would like to define Δ_0 by the formula

$$\Delta_0(a_r, \Gamma(p)) = (\omega_{ij}, \Gamma(p)), \tag{14}$$

where $\{a_r\}$ and $\{\omega_{ij}\}$ are related by equation (11) of §3. However, that equation requires the choice of a fixed reference point p_*, and this choice restricts us to p's such that no edge of $\Gamma(p)$ is collinear with p_*. Assume this restriction for the moment. It is then easy to argue, using 3.3.5, that Δ_0 gives a 1-1 correspondence. Formula (11) shows that this correspondence and its inverse are continuous. That is, subject to the above restriction on p, Δ_0 is a homeomorphism.

We now argue that the restriction on p may be eliminated. First note that the restriction on p is an "open condition"; that is, a specific choice of reference point p_* restricts us to homeomorphic open sets in $PA(\Gamma, p)_0$ and $S(\Gamma, p)$, say $U(p_*)$ and $V(p_*)$, respectively. Next note that by varying p_* in H, the resulting families $\{U(p_*)\}$ and $\{V(p_*)\}$ cover $PA(\Gamma, p)_0$ and $S(\Gamma, p)$, respectively. Finally, note that the value of the right-hand side of (11), is

independent of the choice of p_* (see Appendix A, Corollary A.5). This means that the homeomorphisms $\Delta_0: U(p_*) \to V(p_*)$ fit together compatibly to give a map $\Delta_0: PA(\Gamma,p)_0 \to S(\Gamma,p)$, which is, thus, a local homeomorphism. But, using Corollary 3.3.5 and equation (14) above, it is easy to see that Δ_0 is a 1-1 correspondence. Thus,

$$\Delta_0: PA(\Gamma,p)_0 \to S(\Gamma,p)$$

is a homeomorphism with no restriction needed on p other than that $\Gamma(p)$ be a planar embedding of Γ.

Sometimes a voluntary restriction on p is desirable. For example, we may wish to hold some subset of the vertices fixed, while allowing the rest to vary. One special case of this is important enough to deserve some special mention. This is the case when the fixed vertices comprise the peripheral k-gon Π of Γ. Let $PA(\Gamma,p)_\Pi$ and $S(\Gamma,p)_\Pi$ denote the closed subsets of $PA(\Gamma,p)_0$ and $S(\Gamma,p)$, respectively, in which the peripheral vertices are restricted to certain fixed values. Then, Δ_0 in (14) restricts to a homeomorphism (still denoted Δ_0)

$$PA(\Gamma,p)_\Pi \overset{\to}{\approx} S(\Gamma,p)_\Pi.$$

For simplicity, we usually omit the subscript Π from the notation.

3.5 Remarks on stressed graphs

3.5.1 It is important for the reader to be aware that not every graph $\Gamma(p)$ admits an equilibrium stress. For example, the (solid-line) graph in Figure 3(a) does admit such a stress; the (solid-line) graph in Figure 3(b) does not.

Figure 3. Examples of embedded graphs with and without equilibrium stresses.

Thus, whereas it is relatively easy to deform $\Gamma(p)$ (say, from configuration (a) to (b) above), it may not be possible to deform an associated equilibrium stress compatibly. It is precisely this phenomenon that we encounter when we attempt to deform polyhedra, only now it is expressed in terms of stressed graphs. In Section 4 we show that certain deformations of convex stressed graphs are possible, and this will be sufficient for our applications.

3.5.2 Consider an arbitrary compact, convex polyhedron K with face f_0. By applying a suitable affine transformation A, we may arrange things so that $A(f_0)$ lies in the hyperplane $H = \{(x,y,z) \mid z = 1\}$. We then choose a point p in $\mathbf{R}^3$ outside the polyhedron $A(K)$ but near one of the interior points of $A(f_0)$. By performing a central projection from p, the entire boundary of $A(K)$ can be mapped into $A(f_0)$, provided p is close enough to $A(f_0)$ (see Figure 4). This is known as a Schlegel projection, and the projected 1-skeleton of $A(K)$ is known as a Schlegel diagram. (This proves that such 1-skeleta are planar.) Now, by applying a projective transformation T that fixes H and sends p to the point at infinity corresponding to the direction of the z-axis, we may arrange things so that $TA(K)$ sits over its face $TA(f_0)$ in H. Moreover the Schlegel projection is then replaced by the usual orthogonal projection onto H. Thus, by applying a suitable projective map, in this case TA, we can transform any compact, convex polyhedron K and face f_0 into the context analyzed in this section. We shall call such a transformation (or change of coordinates) a *normalization* of the situation. In our applications, we usually deal only with the normalized objects and leave it to the reader to perform the normalization or its inverse.

Figure 4. Schlegel projection

3.6 Convex polyhedra and convex stressed graphs

The following result is well known to experts, but a proof does not seem to be readily accessible in the literature. We present a proof at the end of Appendix A (Lemma A.9).

3.6.1 Theorem: Suppose $(\{\omega_{ij}\}, \Gamma(p))$ is a stressed graph in H and Σ is the corresponding polyhedral surface sitting over $\Gamma(p)$. Then Σ is the boundary of a convex polyhedron if and only if $\Gamma(p)$ and $\{\omega_{ij}\}$ are convex.

§ 4. Intersecting a convex polyhedron with a half space

In this section we present an algorithm for intersecting a convex polyhedron with a half space. The algorithm has the property that all conversions of numerical data to logical data are independent and consistent. Moreover, assuming a suitable degree of precision in the input data and in the arithmetic that is used, it satisfies the definition of correctness of §1.3. The algorithm is described in §4.2 and a proof of correctness is given in §4.4.

§4.1 Representations and Accuracy

A representation R of a convex polyhedron K is a data structure containing incidence information about K as well as face equations approximating those of K. We may call K a *model* of R. The incidence information, given as logical data, lists the incidences among the vertices, edges, and faces of K, what we call the *face-topology* of K. The face equations, given as algebraic and numerical data, correspond bijectively to the listed faces, with no *a priori* assumptions about their degree of accuracy. Thus, a given R may have many models with nothing in common but their face-topology.

Topological criteria due to Tutte [T] can be used to determine whether a candidate for a representation actually is a representation of some convex polyhedron. Specifically Tutte describes when a given graph admits a convex embedding into the plane with a given sub-k-gon as periphery and possessing a restricted, convex equilibrium stress. Choose the periphery

to be triangular, if possible. Then the stress may be modified to become a convex equilibrium stress (cf. §2.3), and the Maxwell correspondence of §3 implies that the graph is the 1-skeleton of a convex polyhedron. If the periphery cannot be chosen to be triangular, the graph must have a degree-three vertex. Introduce a triangle in place of this vertex, apply the foregoing argument to the modified graph, and then restore the original vertex.

Of course, without some further assumptions on the relationship between the incidence information in R and the face equations of R, the notion of representation is not very useful. For example, if K_0 denotes the polyhedron actually determined by the face equations, K_0 may be the empty set, or it may be some unbounded cone in $\mathbf{R}^3$, both totally unrelated to K. Minimally, we are entitled to require that K_0 satisfy the following:

(1) K_0 is a (bounded) convex polyhedron with exactly one face for each face equation in R. Two such faces meet only if they are listed as incident in R.

(2) Each edge listed in R actually appears in K_0 as the intersection of the appropriate faces. Moreover, the cyclic order of such edges around a given face occurs in K_0 as indicated in R.

It is easy to see that if the face equations of R are sufficiently close to those of K, then K_0 does satisfy (1) and (2).

Note that generally K_0 will have edges other than those described in (2). In fact, each vertex v of R (or, equivalently, of K) determines a subgraph $T(v)$ of the 1-skeleton of K_0. Each edge of $T(v)$ is the intersection of two faces whose counterparts in R are not listed as adjacent to each other but are listed as incident to v. It is not hard to deduce from (1) and (2) that the graphs $T(v)$ are nonempty, connected, and contain no cycles. Thus, each $T(v)$ is either a single vertex or a tree. We call it a *vertex-tree*. By (1), every edge of K_0 belongs to a vertex-tree or arises as in (2).

The features described in (1) and (2) are stable under small perturbation of face equations. However, in general the vertex information of K, carried by R, is not preserved in K_0

because vertices of degree greater than three are unstable with respect to such perturbation. It is precisely this instability phenomenon that has led to most of the considerations of this paper.

Henceforth we consider only representations satisfying the two assumptions above. If we wish to emphasize that they are in force, we may call the representation in question a *good representation*.

The algorithm that we present in §4.2 will require some additional considerations which are more quantitative in nature. In general, the vertex trees $T(v)$ of K_0 can be determined exactly from R only if infinite-precision arithmetic is used. In the event that finite-precision arithmetic is used, the location of each vertex of $T(v)$ may be determined only up to some small radius. Let $N(v)$ denote the set of all points in $\mathbf{R}^3$ whose distance from $T(v)$ is less than or equal to this radius. To include the infinite-precision case, we allow the radius to be zero.

Now choose a positive real number α. Then we call R an α-representation of K if R is a good representation and if the following three conditions are satisfied:

(3) For each v in K, $N(v)$ is contained in the open ball $B(v, \alpha)$ of radius α around v. We shall call $B(v, \alpha)$ a vertex ball belonging to v.

(4) No two such balls meet, nor is any vertex ball with α of the convex hull of any two others.

(5) If a plane meets three or more vertex balls belonging to vertices of a face, then the only vertex balls it meets belong to vertices of that face.

Condition (3) gives a measure of the accuracy of the representation with respect to the model K. It asserts that the total imprecision coming from both the face-equation data of R and the finite-precision computations that may be used is less than α. Condition (4) asserts that the edges and faces of K are large relative to α (or inversely, that α is small relative to these features). It also implies that the angles between consecutive edges in a face are

bounded away from π or 0 and that the dihedral angle between a face and a plane passing near three of its vertices is bounded away from $\pi/2$. Condition (5) relates α to the dihedral angles of K. It will be violated if and only if these angles are too close to 0 or π relative to α.

Given K, conditions (4) and (5) may be formulated as upper bounds on α. Thus, for each convex polyhedron K, there exists a positive number β such that K has α-representations for all $\alpha < \beta$. If R is a representation of K and the total imprecision is known *a priori*, then the upper bounds of (4) and (5) may be estimated from R. When both estimates are greater than the total imprecision, then, for any choice of α in between, R is an α-representation. If one of the estimates is less than the total imprecision, then we must reject the representation as being insufficiently precise.

§4.2 The algorithm

Let α be a small positive number. The algorithm takes as input an α-representation R of a convex polyhedron K and a linear inequality $\lambda \geq 0$ defining a half space. Its output is a representation R^* of a convex polyhedron K^*. We may call R^* the *intersection-representation*.

The algorithm proceeds as follows. Let P be the plane $\lambda = 0$. For each vertex v of K, the algorithm attempts to determine by means of a numerical computation the side of P on which v lies. Such a determination will succeed exactly when the computed location of v has distance from P greater than or equal to α. In this case, the algorithm reports the side containing the computed location. If the attempted determination fails, that is, if the computed location is strictly within α of P, then v is reported to be on P.

If none of the vertices of K are reported on the side $\lambda > 0$, then the algorithm returns the null representation. Similarly if none are on $\lambda < 0$, then the algorithm returns R.

Suppose that three or more vertices of a face F are reported on P, and some vertices of K are reported on each side of P. By property (5) of an α-representation, this can occur for at most one face F of K. In this case, the algorithm computes the inward pointing normals of F

and P (i.e., for F this means pointing inward into K, for P this means pointing inward into $\lambda \geq 0$), and then it computes the dot product of these normals. It returns the null representation if the dot product is negative and R if it is positive. Note that, by property (4) of an α-representation, the absolute value of this dot product is bounded below. Therefore, with only a modest assumption on total precision, that is, on the input α, the algorithm is able to make the decision in this case.

So let us now assume that each face has at most two vertices reported on P and that some vertices of K are reported on each side of P. To construct the incidence data for the intersection-representation in this case, partition each edge of K that intersects P. This is done solely as a logical computation based on the location of endpoints of the edge relative to P. Discard segments on the "wrong" side of P, and to the remaining set of new edges, adjoin all the edges of K both of whose endpoints are reported to be on P or on the correct side of P. Finally, determine new edges by joining appropriate partition points constructed above, or a partition point with a vertex reported on P, as the case may be. This gives all the edges and vertices for the intersection-representation. The faces are of three kinds: the face determined by P itself, portions of faces of K which are divided by P, and faces of K lying on the "correct" side of P. Again, these are all determined by logical computation, as are the required incidences. For the face determined by P, we have the face equation $\lambda = 0$. For the remaining faces, the face equations come directly from R.

This concludes our description of the algorithm and of the output representation R^*. We need to show, however, that R^* actually is a representation.

4.2.1 **Theorem:** *R^* is a good representation.*

Proof: The null representation and R itself are both good representations, so we may assume that the algorithm has reported vertices on both sides of P, at most two of which are on a given face.

Recall that K_0 is the bounded, convex polyhedron given by the face equations listed in R. Let K^*_0 be the analogous convex polyhedron for R^*. To show that R^* is a representation, we refer the reader to the second paragraph of §4.1 and the result of Tutte mentioned there. Tutte's result asserts that it is sufficient to show that the edge-graph G^* given in R^* is 3-connected and planar. Let G^*_0 denote the edge-graph of K^*_0; by Tutte [T], together with the Maxwell correspondence of Section 3, G^*_0 has these properties. Moreover, G^* may be obtained from G^*_0 by collapsing each vertex tree to a point . It is not hard to show that collapsing an edge of a vertex tree does not destroy planarity (e.g., the reverse operation cannot eliminate Kuratowski subgraphs). Thus G^* is planar. To prove that it is 3-connected, observe that the link of every vertex-tree is a simple closed curve. (By the link of a subset S of G^*_0, here, we mean the union of all edges e of G^*_0 such that e is disjoint from S but e belongs to a face meeting S.) Collapse a single tree to a point p; this does not affect the link. If p and another point are selected in the resulting graph, the 3-connectivity of G^*_0 implies that each component of the complement of these points meets the link of p. Since one point cannot disconnect a simple closed curve, this proves that the complement is connected. This same conclusion follows even more easily for the complement of two points each different from p. Thus, the collapse of a single tree in G^*_0 results in a 3-connected planar graph. Now continue this process, collapsing all the trees in G^*_0.

It remains to show that R satisfies conditions (1) and (2) that define a good representation. We first observe that K^*_0 is bounded. To see this, consider the subset M of K_0 on which the linear function λ assumes its maximal value. M is either a single point, a line or a single face of K_0. In any case, M, P, and the planes of faces of K_0 incident to M enclose a bounded region that contains K^*_0. Clearly every face equation listed in R^* gives a face of K^*_0, because such an equation either corresponds to a face or portion of a face of K_0 (since R is good), or it corresponds to the face determined by P. By construction, K^*_0 has no additional faces. If two faces of K^*_0 meet in an edge, then by the definition of R^* and the fact that R is good, these faces are listed as incident in R^*. Thus R^* and K^*_0 satisfy property (1)

of good representations. For property (2), consider the edges listed in R^* that correspond to edges of K, or portions of edges of K, but do not have both endpoints reported on P. Such edges occur in K_0 as intersections of appropriate faces and hence occur in K^*_0. The remaining edges of R^* are portions of intersections of face planes of K_0 with P. That these edges all occur in the appropriate order in the boundaries of faces follows from the corresponding fact for K_0 and the construction of R^*. Thus, R^* and K^*_0 satisfy (2). □

By Theorem 4.2.1, the algorithm produces a valid, consistent geometric object. It remains to demonstrate that our algorithm is correct for the operation of intersection, as per the definition of §1.3. In §4.4, *we show that this is the case, provided that* α *is sufficiently small.* It is also desirable to improve Theorem 4.2.1 so that we know that R^* is not merely a good representation but also that it is a β-representation, where β is suitably related to α. We also do this in §4.4 again under the assumption that α is sufficiently small. In §4.3 we prove some of the technical results necessary for these demonstrations.

4.3 The main technical results

In this section, we present a technical result, Theorem 4.3.3, which we need for our demonstration that the algorithm of §4.2 is correct. Actually, the demonstration requires a similar but more general result, which, because its proof is substantially more complicated, we present in Appendix C.

We continue to use the terminology and notation of §4.2. Thus, in particular, K denotes the convex polyhedron of the algorithm, and λ denotes the given linear form. We cannot move vertices of K arbitrarily, even by small amounts, and expect to find another polyhedron with those vertex locations and the same face-topology as K. However, Theorem 4.3.3 tells us that we can make a small arbitrary vertex change in the "λ-direction" without disturbing the face-topology of K, provided we allow an accompanying change in other directions.

The proof of Theorem 4.3.3 deals entirely with a stressed equilibrium graph for K. Thus, we suppose that we have normalized the situation as discussed in §3.5.2. In particular, K sits

over a convex 3-connected graph $\Gamma(p)$ in the plane $H = \{(x,y,z) \mid z = 1\}$, and P is the plane $X = 0$. We suppose, as usual, that the vertices p_i of the graph have been numbered so that $p_1, p_2, \cdots, p_k$ are peripheral and the rest interior. Call an edge of $\Gamma(p)$ *vertical* if its endpoints have the same x-coordinate. The proof of Theorem 4.3.3 makes use of the fact that every such graph contains a spanning tree with no vertical edges. We prove a slightly stronger result in Lemma 4.3.2, but first we show how such a spanning tree is used.

4.3.1 Lemma: Let $\Gamma(p)$ be as above and suppose that it has a spanning tree T with no vertical edges. For each edge e_{ij} *not* in T, let ω_{ij} be an arbitrary stress constant. Then there exists a unique assignment of stress constants to the edges of T, hence a stress $\{\omega_{ij}\}$ on Γ, such that the x-equilibrium equation

$$\sum_{j=1}^{n} \omega_{ij}(x_i - x_j) = 0$$

is satisfied for each $i = 1, 2, \ldots, n$. Here, x_i denotes the x-coordinate of p_i.

Proof: Suppose that p_1 is a leaf of T and e_{12} is an edge of T. All the stress constants ω_{1j}, $j > 2$, are already determined, and $x_1 - x_2 \neq 0$. Thus, the above equation, for $i = 1$, can be used to solve for ω_{12}. The remaining stress constants are found similarly starting at the leaves of T and working towards the root. Then, x-equilibrium is guaranteed by construction at every vertex except possibly at the root of T. But, in fact, x-equilibrium holds there as well, because the sum of all the expressions

$$\sum_{j=1}^{n} \omega_{ij}(x_i - x_j)$$

$i = 1, 2, \ldots, n$, is formally equal to zero. □

Next we show that not only does $\Gamma(p)$ contain a spanning tree with no vertical edges but it also contains one with no edges that are near-vertical. We say that an edge is *near − vertical* if its endpoints have x-coordinates that differ by less than some prescribed,

small, positive quantity ν. The precise choice of ν is not important. We require only that no face of $\Gamma(p)$ have consecutive edges that are near-vertical. Note that, by the convexity of K and the fact that it 'sits over' $\Gamma(p)$, there exist sufficiently small ν such that this property is satisfied.

4.3.2 Lemma: Let $\Gamma(p)$ be as above, and assume that at most one of its peripheral edges is near-vertical. Then $\Gamma(p)$ has a spanning tree with no near-vertical edges.

Proof: We need only show that the deletion of all near-vertical edges leaves $\Gamma(p)$ connected. Since at most one peripheral edge is near-vertical, the modified periphery remains connected. Since no vertex is incident to a pair of consecutive near-vertical edges and $\Gamma(p)$ is convex and 3-connected, each interior vertex is incident to some non-near-vertical edge going to the left. Thus, each interior vertex may be connected to the modified periphery by a leftward moving path consisting of non-near-vertical edges. □

Note that when the periphery of $\Gamma(p)$ is triangular, then it automatically satisfies the hypothesis of Lemma 4.3.2.

4.3.3 Theorem: Let K be as above, and let X_0 denote the n-tuple of x-coordinates of the vertices of K. There exists an $\delta > 0$ such that if X_1 is any n-tuple satisfying

$$|X_0 - X_1| < \delta,$$

then there is a convex polyhedron K_1 with the same face-topology as K and having X_1 as the n-tuple of x-coordinates of its vertices. Furthermore, K_1 may be chosen so that its vertices vary smoothly as a function of X_1.

Proof: Let $(\{\omega_{ij}\}, \Gamma(p))$ be the convex stressed graph in H corresponding to K. For simplicity, we deal here only with the case in which the periphery of $\Gamma(p)$ is a triangle. A more general situation is considered in Appendix C.

Set $X_0 = (x_1, x_2, \ldots, x_n)$ and $X_1 = (x'_1, x'_2, \ldots, x'_n)$. Since the periphery of $\Gamma(p)$ is non-degenerate, the same will be true when x_1, x_2, x_3 are replaced by x'_1, x'_2, x'_3, respectively, provided $|X_0 - X_1|$ is suitably small. Our first restriction on X_1, then, is that it satisfy $|X_0 - X_1| < \delta_1$, for such a suitably small δ_1.

Next, choose a spanning tree T with no near-vertical edges, and let $\delta_2 = \min\{|x_i - x_j| : e_{ij} \text{ is an edge of } T\}$. Our second restriction on X_1 is that $|X_0 - X_1| < \delta_2$ 2. This insures that no edge of T becomes vertical when the x-coordinates of its endpoints are replaced by the corresponding coordinates of X_1.

We then define a new stress $\{\omega'_{ij}\}$ for Γ as follows. For every edge e_{ij} not in T, set $\omega'_{ij} = \omega_{ij}$. For the remaining edges, proceed as in Lemma 4.3.1, using T and the coordinates of X_1. Clearly the resulting stress is a smooth function of X_1. Since $\{\omega'_{ij}\} = \{\omega_{ij}\}$ when $X_1 = X_0$, and $\{\omega_{ij}\}$ is a convex stress (i.e., ω_{ij} is negative for e_{ij} peripheral and positive otherwise), there exists a positive δ_3 such that $\{\omega'_{ij}\}$ is a convex stress for all X_1 satisfying $|X_0 - X_1| < \delta_3$. Our final restriction on X_1, then, is that it satisfy this last inequality.

In summary, then, the positive quantity δ of the theorem may be defined as $\min\{\delta_1, \delta_2, \delta_3\}$.

Now set $p'_r = (x'_r, y_r, 1) \in H$, $r = 1,2,3$, where y_r is the y-coordinate of p_r. By the choice of δ, these are the vertices of a non-degenerate triangle in H. Since $\{\omega'_{ij}\}$ is a convex stress on Γ, we may apply Theorem 2.1.3 to obtain unique points $p'_r, r > 3$, satisfying: (1) $\{\omega'_{ij}\}$ is a restricted convex equilibrium stress for $\Gamma(p')$; (2) the p'_r depend smoothly on the ω'_{ij}. $i > 3$ or $j > 3$ (hence, smoothly on X_1). We claim that X_1 is precisely the n-tuple of x-coordinates of p'. For, by construction, the coordinates of X_1 satisfy x-equilibrium with respect to $\{\omega'_{ij}\}$. In matrix form, this is just the first of the two equations (8) of §2, or, equivalently, the first of the two equations (9). But, by Corollary 2.2.2, this same equation gives the x-coordinates of $p'_4, p'_5, \ldots, p'_n$ in terms of the x-coordinates of p'_1, p'_2, p'_3, which coincide with x'_1, x'_2, x'_3, respectively, by definition. Thus the claim is verified.

Now modify the stress constants ω'_{ij}, $i,j < 4$, as described in Theorem 2.3.1, so as to obtain a convex equilibrium stress, still denoted $\{\omega'_{ij}\}$, on $\Gamma(p')$. By (2) above, and by assertion (2) of Theorem 2.3.1, the stressed graph $(\{\omega'_{ij}\}, \Gamma(p'))$ depends smoothly on X_1.

The results of §3 then produce a convex polyhedron K_1 corresponding to $(\{\omega'_{ij}\}, \Gamma(p'))$. The last assertion of the previous paragraph, together with the results of §3, shows that the vertices of K_1 depend smoothly on X_1. Moreover, they have the correct x-coordinates, by construction. □

4.3.4 **Remark:** Since the vertices of K_1 depend smoothly on X_1, it is clear that we may choose K_1 to be close to K by choosing X_1 suitably close to X_0. More precisely, for every $\varepsilon > 0$, there is a $\delta > 0$ such that if $|X_0 - X_1| < \delta$, then vertices of K_1 are within ε, of the corresponding vertices of K. To compute such a δ in terms of ε, one must factor the smooth map $X_1 \dashrightarrow K_1$ into its three component parts, $X_1 \dashrightarrow \{\omega'_{ij}\}$, $\{\omega'_{ij}\} \dashrightarrow p'$, and $(\{\omega'_{ij}\}, \Gamma(p')) \dashrightarrow K_1$, and then estimate moduli of continuity for these. For our purposes here, it is sufficient to note that, given K and ε, these steps can be carried out. In addition, one can factor in the effect of normalization and its inverse. A similar but more complicated analysis applies to the more general method used in Appendix C. In §4.4 we discuss how δ may be estimated when only R is known.

4.4 Correctness and Accuracy

To establish the correctness of our intersection algorithm according to the definition of §1.3, it is sufficient to show that a model for the input α-representation R exists whose actual intersection with the half-space $\lambda \geq 0$ is a model for the output representation R^*. Note that we are not imposing any *a priori* standard of accuracy on R^*. Of course, we could do so by requiring that R^* be a $\beta-$*representation* of the actual intersection, for a prescribed $\beta > 0$. We touch on this below. In any case, we are left with the problem of producing a suitable model for the input representation. In those cases in which the algorithm produces the null

representation or R itself, this problem is easily solved. For in those cases, either a small translation will perturb K so that it no longer meets the plane P or small rotation followed by translation has this effect. In either case, this perturbation of K is then the desired model of R.

For the generic case when vertices are reported on both sides of P and no face has more than two vertices reported on P, we again start with the given polyhedron K and attempt to perturb it slightly so as to meet our requirements. That is, we attempt to perturb K, without changing its face-topology, so that precisely those vertices that are reported on P are actually moved on to P and no vertex moves from one side of P to the other. If we can do this, then the perturbed image of K is again clearly the desired input model.

Let us say that a face is *reported* to *meet* P if it has vertices reported on both sides of P. Suppose first that a triangular face F is reported to meet P. We then normalize the situation as in §4.3, projecting K onto F, moving F to the plane H and P to the plane $x = 0$. The projected image of P in H is a line L that divides the graph $\Gamma(p)$. Vertices within α of P are projected to vertices in $\Gamma(p)$ within α of L. We attempt to move these vertices onto L; that is, to change their x-coordinates to 0. By Theorem 4.3.3, this will be possible if α is less than a certain quantity δ.

A similar argument may be used if P fails to meet a triangular face as above but is suitably close to an order-three vertex. For then we need only slice off this vertex in a direction transversal to P, then apply the preceding argument, and finally restore the order-three vertex.

We next consider the remaining possibilities. It is possible, for example, that there are faces reported to meet P none of which is triangular and P is not close to an order-three vertex. In this case, the same discussion applies as in the previous paragraph, only we use the more general result of Appendix C instead of Theorem 4.3.3. The same conclusion holds. Finally, if no face is reported to meet P and no order-three vertex is close P, but there are ver-

tices

reported on both sides of P, then K must be divided onto two pieces by a simple closed curve consisting of edges that lie on P or nearly lie on P (i.e., each endpoint of such an edge is reported to be on P). In this case, we choose such an edge and perturb K slightly so that the edge is moved onto P. Then we choose either of the two faces f_1 and f_2 containing the edge and project K onto the plane of this face, say f_1, using a central projection from a point close to f_1, not in K but on the plane P (as usual). Normalize as before so that the plane of f_1 becomes H, P becomes $x = 0$, and the projection is orthogonal. We get a stressed graph $\Gamma(p)$ as before, only its periphery now encloses the image of both f_1 and f_2 which are exterior to one another. The faces are all convex, the interior stress constants are positive, and the peripheral ones are negative, just as before. This case can be handled just as the previous one and is discussed in more detail in Appendix C.

Thus, in all cases, we can obtain the desired deformation provided α is small enough. We now formulate this as a theorem.

4.4.1 Theorem: Given any convex polyhedron K and linear form λ, there is a $\delta > 0$ such that the intersection algorithm of §4.2 is correct for all α-representations R of K satisfying $\alpha < \delta$. Moreover, δ may be chosen to be independent of λ.

Proof: Only the second assertion still requires verification. Without loss of generality, we restrict attention to forms $\lambda = ax + by + cz + d$, such that $a^2 + b^2 + c^2 = 1$. Those λ for which the output of the algorithm is the null representation or R require no δ: that is, no restriction on α is needed for those branches of the algorithm. The remaining λ satisfy the condition that vertices are reported on both sides of $\lambda = 0$ but at most two from each face are reported on $\lambda = 0$. This is a closed, bounded condition on the coefficients of λ, and so this set of λ comprise a compact set C.

Choose a particular λ_1 in C, and let δ_1 be the value of δ given by the discussion above, i.e., by Theorem 4.3.3 after normalizing the situation. Let us say that a value of δ is *suitable* for λ if, whenever $\alpha < \delta$, K may be perturbed so that precisely the vertices within α of $\lambda = 0$ are moved onto that plane and no vertex changes sides. We claim that $\delta_1/2$ is a suitable value of δ for all λ in some neighborhood of λ_1 in C. To verify this claim, choose an $\alpha < \delta_1/2$ and select those vertices of K within α of $\lambda = 0$. Let $T(\lambda)$ denote the projective transformation that we use to normalize K. For λ close to λ_1 , $T(\lambda)$ may be chosen to be a smooth function of λ, and hence $T(\lambda)$ is close to $T(\lambda_1)$. Because the x-coordinates of the selected vertices of $T(\lambda)(K)$ are less than α, it then follows that the x-coordinates of the corresponding vertices of $T(\lambda_1)(K)$ are less than δ_1. Therefore, as in Theorem 4.3.3, there exists a perturbation h of $T(\lambda_1)(K)$ sending the selected vertices onto $x = 0$, etc. Now when λ is close to λ_1, there is an affine transformation A of R^3 that preserves orientation and transforms λ_1 to λ: $\lambda = \lambda_1 A$. Define a mapping of g of R^3 by $g = A^{-1} T(\lambda_1)^{-1} hT(\lambda_1)$, and consider any of the selected vertices p_i of K. Since $T(\lambda_1)(p_i)$ is taken onto $x = \hat{0}$ by h, and $T(\lambda_1)$ takes $\lambda_1 = 0$ to $x = 0$, $\lambda g(p_i) = 0$, as desired. Similarly g preserves the half-spaces determined by $\lambda = 0$. Thus, $g(K)$ is a perturbation of K with desired properties, which verifies the claim.

Since finitely many such neighborhoods cover C, each with a suitable value of δ, the smallest of these values is suitable for all λ in C, and therefore, for all λ. □

We now turn to the question of the accuracy of the representation produced by the algorithm. There are (at least) two senses in which the term accuracy can be used here. First, we can ask how closely the output representation R^* represents some model. Secondly, we can ask how closely such a model approximates the true intersection of K and $\lambda \geq 0$.

The first question has a relatively straightforward mathematical answer.

4.4.2 Theorem: There exists a positive constant κ depending only on K, such that if R is an α-representation of K and α is less than the constant δ of Theorem 4.4.1, then each vertex of

the model K^* of R^* constructed as in the correctness proof above is within $\kappa\alpha$ of its location computed using R^*

Proof: The proof consists of looking closely at the construction of the modified input model in our demonstration of the correctness of the algorithm. Suppose first that the algorithm reports all vertices on one side of P or on P. If this side corresponds to $\lambda \geq 0$, then K^* equals a small translate of K into this side of P. The length of the translation need be at most 2α, so that $R^* = R$ is a 2α-representation of K^* in this case. If the side reported corresponds to $\lambda \geq 0$, then K^* equals the empty polyhedron.

The case in which the algorithm reports at least three vertices of a face F on P and some vertices on both sides of P is more complicated. In this case, we rotate K through a small angle around the intersection line $P \cap F$ until F lies on P, and then we perform a small translation, say of length α, making the polyhedron disjoint from P. If the result lies in $\lambda \geq 0$ then K^* equals this resulting polyedron. Otherwise, K^* is empty. It remains to show that there is a constant M such that each vertex of K^* is less than $M\alpha$ from the corresponding vertex of K. For then the distance between such a vertex and that computed via $R^*(=R)$ is less than $(M + 1)\alpha$, as desired. But, for any fixed set of three or more vertices of F within α of P and line segment of intersection of P and F, the distance moved by vertices under the rotation is a smooth function of the angle of rotation, and this angle in turn has an upper bound of the form $m\alpha$. Similarly, the locations of the endpoints of the line segment are determined up to some $m'\alpha$. Thus, there is a constant M' such that $M'\alpha$ bounds the distance moved by vertices under rotation. An additional α takes care of the small translation, so that $M = M' + 1$ is the desired constant. Note that we may repeat this argument for any set of three or more vertices on any face, thus finding a single M independent of λ in this case.

Next we choose any λ in the compact set C described in the proof of the previous theorem. Let T be the projective transformation that maps K into its normalized position over the graph $\Gamma(p)$ in H. For any polyhedron K' with n vertices, let $v(K')$ denote the n-tuple of

vertices, and let $x(K)$ be the n-tuple of x-coordinates of these vertices. Let τ be the modulus of continuity of T on some large closed, bounded region. Now apply the construction of Theorem 4.3.3 or Appendix C, as the case may be, to the polyhedron $T(K)$ and some n-tuple of x-coordinates $u = (u_1, u_2, \ldots, u_n)$. The result is a polyhedron K_1 satisfying $x(K_1) = u$ and depending smoothly on u. Thus, for u ranging in a compact ball around $x(T(K))$, there is a constant M such that $|v(T(K)) - v(K_1)| < M|x(T(K)) - u|$. In our perturbation of K to construct a model of R^*, the vertices that are moved onto P move no more than α units in the direction orthogonal to P. Since T is just a rigid motion followed by a central projection from a point close to K but outside of K, the coordinates in $x(T(K))$ of vertices that we want to perturb are all bounded by α. For our construction of the desired model, we set $u = 0$. Thus, we get $|v(T(K)) - v(K_1)| < M\alpha$. Next, we apply the inverse transformation T^{-1} to $T(K)$ and K_1 and conclude that $|v(K) - v(T^{-1}(K_1))| < M\tau^{-1}\alpha$. Since R is an α-representation of K, we see that a vertex computed using R is within $(M\tau^{-1} + 1)\alpha$ of the corresponding vertex in $T^{-1}(K_1)$. Finally, recall that to construct the model for the output R^*, we may also have to introduce new vertices where edges of $T^{-1}(K_1)$ cut the plane P. Such an edge has endpoints at least $\rho\alpha$ units from P, for some ρ that depends only on T, and so the angle it forms with P is bounded below. It follows that the indeterminacy for such vertices is bounded by $N\alpha$ for some constant N. Thus, setting K^* equal to the actual intersection of $\lambda \geq 0$ with $T^{-1}(K_1)$ and $\kappa = \max\{N, M\tau^{-1} + 1\}$, we see that the vertices of K^* are within $\kappa\alpha$ of those computed using R^*.

We can now combine the foregoing argument with a compactness argument similar to that of Theorem 4.1.1 to conclude that κ may be chosen independently of λ. We shall omit these details. □

In general, $\kappa\alpha$ could be too large to allow us to conclude that R^* is a $\kappa\alpha$-representation of K^*. However, for small α this will be the case. In fact, it is again possible to find a δ_1 independent of λ such that R^* is a $\kappa\alpha$-representation of K^* whenever $\alpha < \delta_1$.

The question of how closely the model K^* approximates the actual intersection of K with $\lambda \geq 0$ is not so straightforward as the first question because we have not established a standard for comparing polyhedra with different face-topologies.

One final question concerning accuracy remains to be addressed. Namely, given an α-representation of R of a convex polyhedron K and linear form λ, how do we estimate suitable constants δ and κ (relative to the given λ) as in Theorems 4.4.1 and 4.4.2? The answer is to mimic the constructions in these theorems (and in Theorem 4.3.3 and Appendix C) using the data of R. The constructions are all smooth functions of the numerical data so that, in the end, the quantities δ' and κ' obtained will deviate from δ and κ by amounts not greater than $M\alpha$, where M is a constant depending only on the face-topology and, perhaps, some gross boundedness assumptions. As long as $\delta' > (M+1)\alpha$, we may use $\delta' - M\alpha$ to estimate δ. Similarly, we may use $\kappa' + M\alpha$ to estimate κ. If $\delta' \leq (M+1)\alpha$, then we must reject R as being insufficiently precise.

4.5 Iteration

The intersection algorithm of §4.2 may be applied to the output representation R^* and some new half-space $\lambda^* \geq 0$ provided that R^* is an α^*-representation of K^* for suitably small α^*. We have seen that this will be the case provided we have been able to choose our initial α small enough. If α is not suitably small, R^* will still satisfy conditions (1) - (3) of §4.1, but conditions (4) and (5) may fail. This means, essentially, that K^* has certain features that are too small. One option, then, is to return to the original input representation R and improve its accuracy. Another option might be to single out the undesirable small features of K^* and attempt to collapse or "shrink" them to points, much as we collapsed vertex trees in the proof of Theorem 4.2.1. Such a shrinking process can, in fact, be carried out, but it has the effect of moving *all* the vertices of the polyhedron. Thus, once some features are collapsed, others may have been made undesirably small. Moreover, the accuracy of the modified represen-

tation may not be as good as that of R^*.

5. Conclusions

This paper explores the paradigm of structuring code so that all conversions from numerical data to symbolic data are explicit and logically independent. The paradigm is applied successfully to the problem of intersecting convex polyhedral objects, and this leads to an algorithm which is correct in the sense of §1.3. Since correct algorithms are clearly robust with respect to numerical-to-logical conversions, the paper produces a robust algorithm for intersecting convex polyhedra.

The key tool used in the paper is a correspondence between stressed planar graphs and convex polyhedra, which allows manipulation of convex polyhedra by elementary graph-theoretic and algebraic means. This tool is used to deform a convex polyhedron so that certain of its vertices may be placed on a prescribed plane. It may also used to remove certain small features from a convex polyhedron.

Prior to the work of this paper, the notion of attempting to make conversions to symbolic data independent was explored in the implementation of an algorithm for intersecting nonconvex polyhedral objects [HHK2]. Although the algorithm did fail on certain inputs, its numerical robustness even in the nonconvex case was significantly better than existing algorithms. The proof techniques of this paper while limited to convex polyhedra do point out certain places where problems are likely to arise in intersecting nonconvex polyhedra. For example, the underlying graph of a convex polyhedron is always 3-connected, and essential use was made of this property in various proofs. A polyhedron in which a face has a hole (see Fig. 5)

may give rise to a disconnected graph. Similarly if two faces intersect in a disconnected set (see Fig. 6) the underlying graph will not be 3-connected. This suggests that when either of these situations occur a face be subdivided so that the graph will be 3-connected.

The topology of a convex polyhedron places only minimal geometrical constraints on its vertices, such as the obvious constraint that all vertices on a face lie on a plane. In the case of nonconvex polyhedra, however, there is a deep interaction between geometry and topology. In this case, for example, the topology can force three vertices to lie on a straight line. This allows the construction of a topology to force complex geometrical relations between vertices on widely separate faces. These constraints can make questions concerning the incidence of vertices on faces logically dependent, and thus the above paradigm fails. However, these constraints seem to arise only when the intersection of two faces is not a connected region. Thus checking for this situation and partitioning a face when it occurs may improve robustness substantially.

Consider two rectangular bars crossing (see Fig. 7). There is a tendency to think of ABCD as a straight line that is the intersection of F_1 and F_2. However, it may be three line segments. In fact F_2' and F_2 may not lie on the same plane. In Figure 8, F_1' and F_1 lie on the same plane. By making use of this construction we can force vertices to lie on a straight line and hence can force relations between various vertices by means of theorems such as Pascal's Theorem.

Bibliography

[BSST] Brooks, R.L., C.A.B. Smith, A.H.Stone, and W.T.Tutte, *The Dissection of Rectangles into Squares*, Duke Math.J. 7 (1940) 312-340.

[C] Connelly, Robert, *Rigidity and Energy*, Invent. Math. 66 (1982) 11-33.

[CW] Crapo and Whitely, *Statics of Frameworks and Motions of Panel Structures, a Projective Geometric Introduction*, Structural Topology 6 (1982)

[GH] Greenberg, M. and Harper, J.R., "Algebraic Topology: A First Course," Benjamin/Cummings, Menlo Park, CA 1981

[HY] Hocking, J.G. and Young, G.S., "Topology," Addison-Wesley, Reading, MA 1961

[G] Grunbaum, Branko, "Convex polytopes," John Wiley and Sons, New York 1967

[KKH1] Hopcroft, J.E., C.M. Hoffmann, and M.S. Karasick, *Towards Implementing Robust Geometric Computations*, Proc. of the Fourth Annual Symposium on Computational Geometry, Urbana-Champaign, June 6-8, 1988 106-117

[KKH2] Hopcroft, J. E., C. M. Hoffmann, and M. S. Karasick, *Robust Set Operations on Polyhedral Solids*, IEEE Computer Graphics, to appear.

[K] Kahn, P.J., *Geometric Applications of Cosheaf Homology*, in preparation

[K2] Kahn, P. J., *Perturbing Non-Generic Convex Polyhedra*, in preparation

[M] Maxwell, James C., "On Reciprocal Figures and Diagrams of Forces," Phil. Mag. Series 4, 27 (1864), 250-261

[T] Tutte, W.T., *How to Draw a Graph*, Proc. London Math. Soc. (3) 13 (1963) 743-768

[W] Whiteley, Walter *Motions and Stresses of Projected Polyhedra*, Str ctural Topology 7 (1982) 13-38

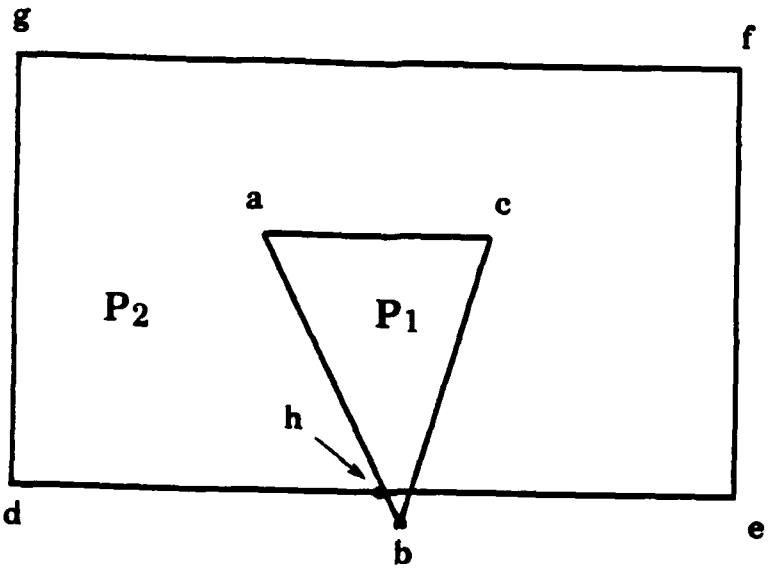

Figure 1: Intersecting two polygons

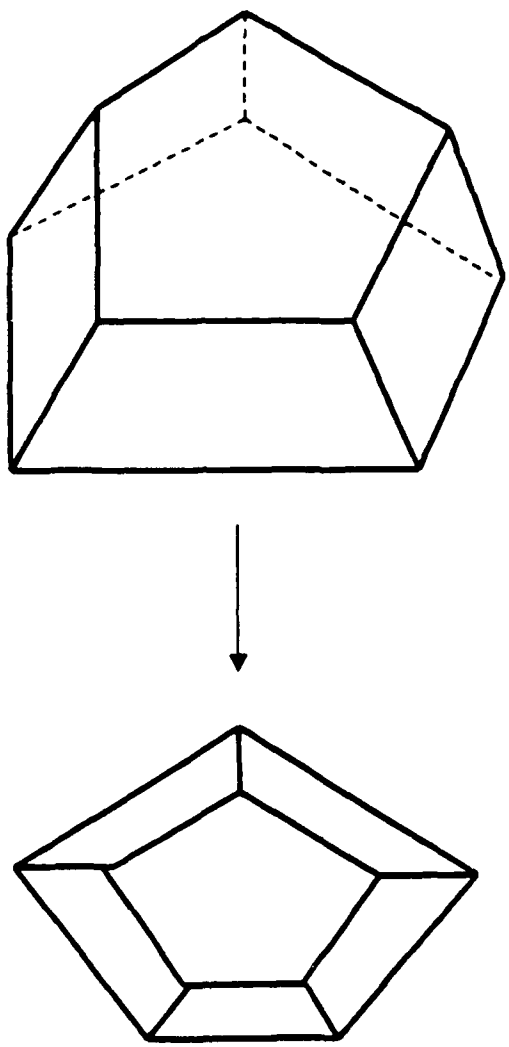

Figure 2: Polyhedron sitting over plane and the graph of its projection

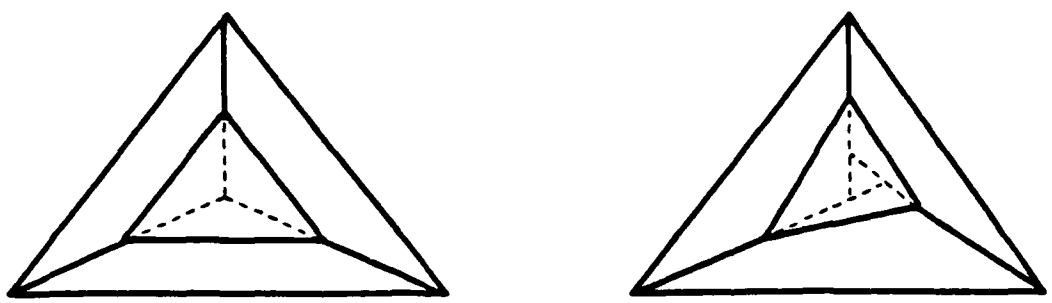

(a) (b)

Figure 3: Examples of embedded graphs, with and without equilibrium stresses.

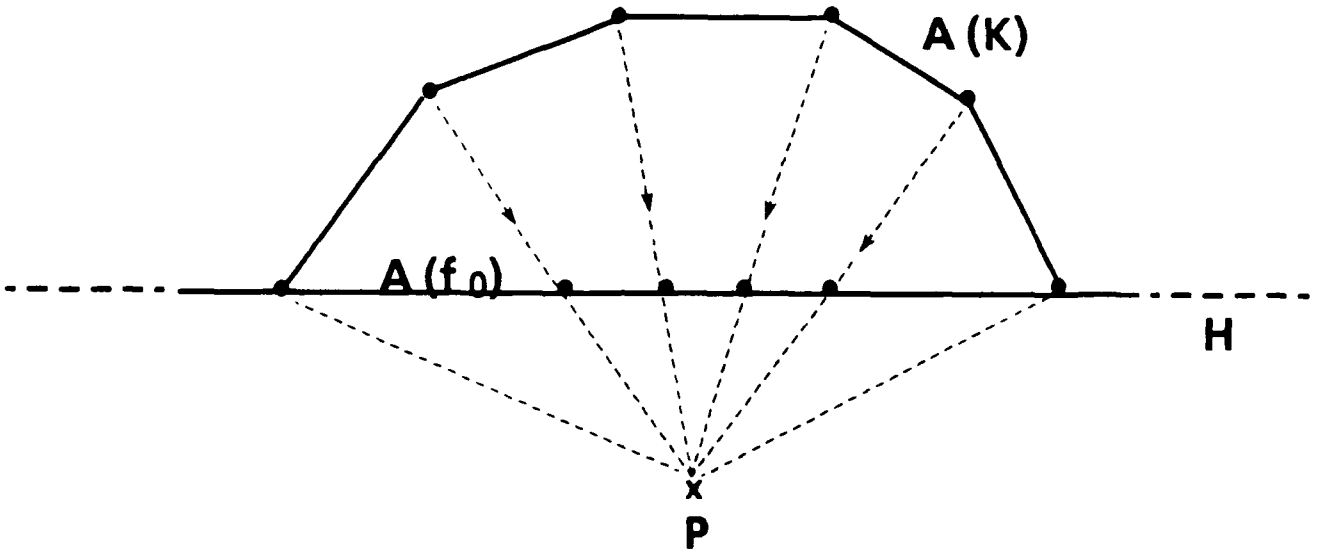

Figure 4. Cross-section of a Schlegel projection

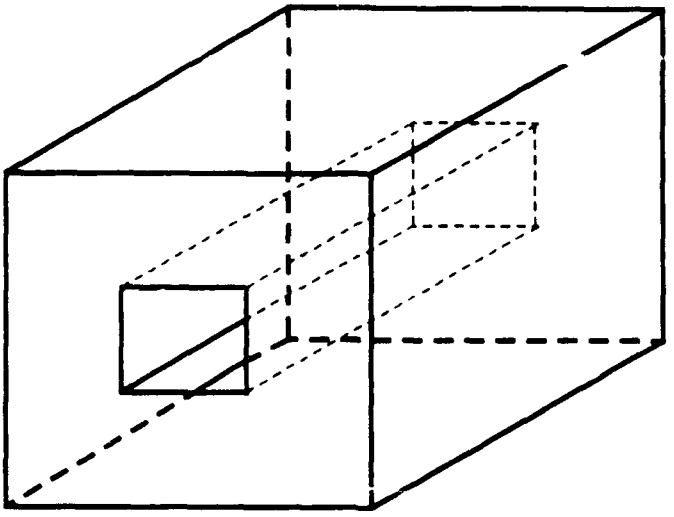

Figure 5. Solid whose graph is disconnected

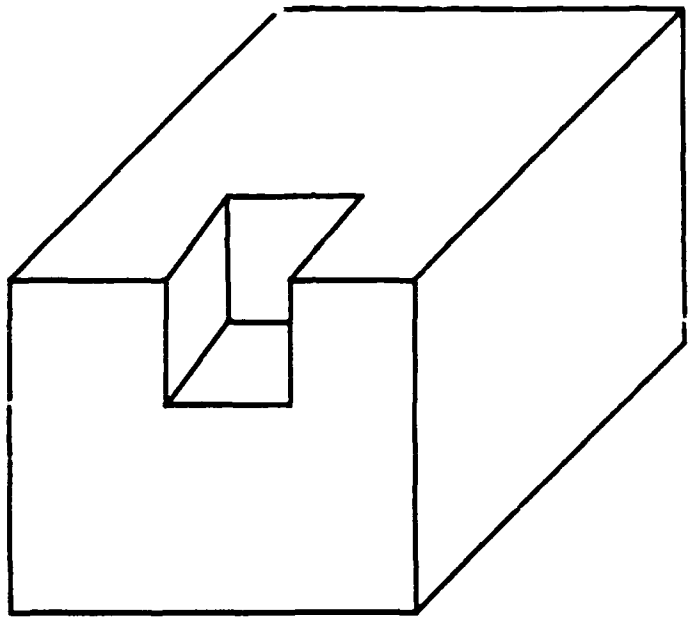

Figure 6: Two faces that meet in a nonconnected set.

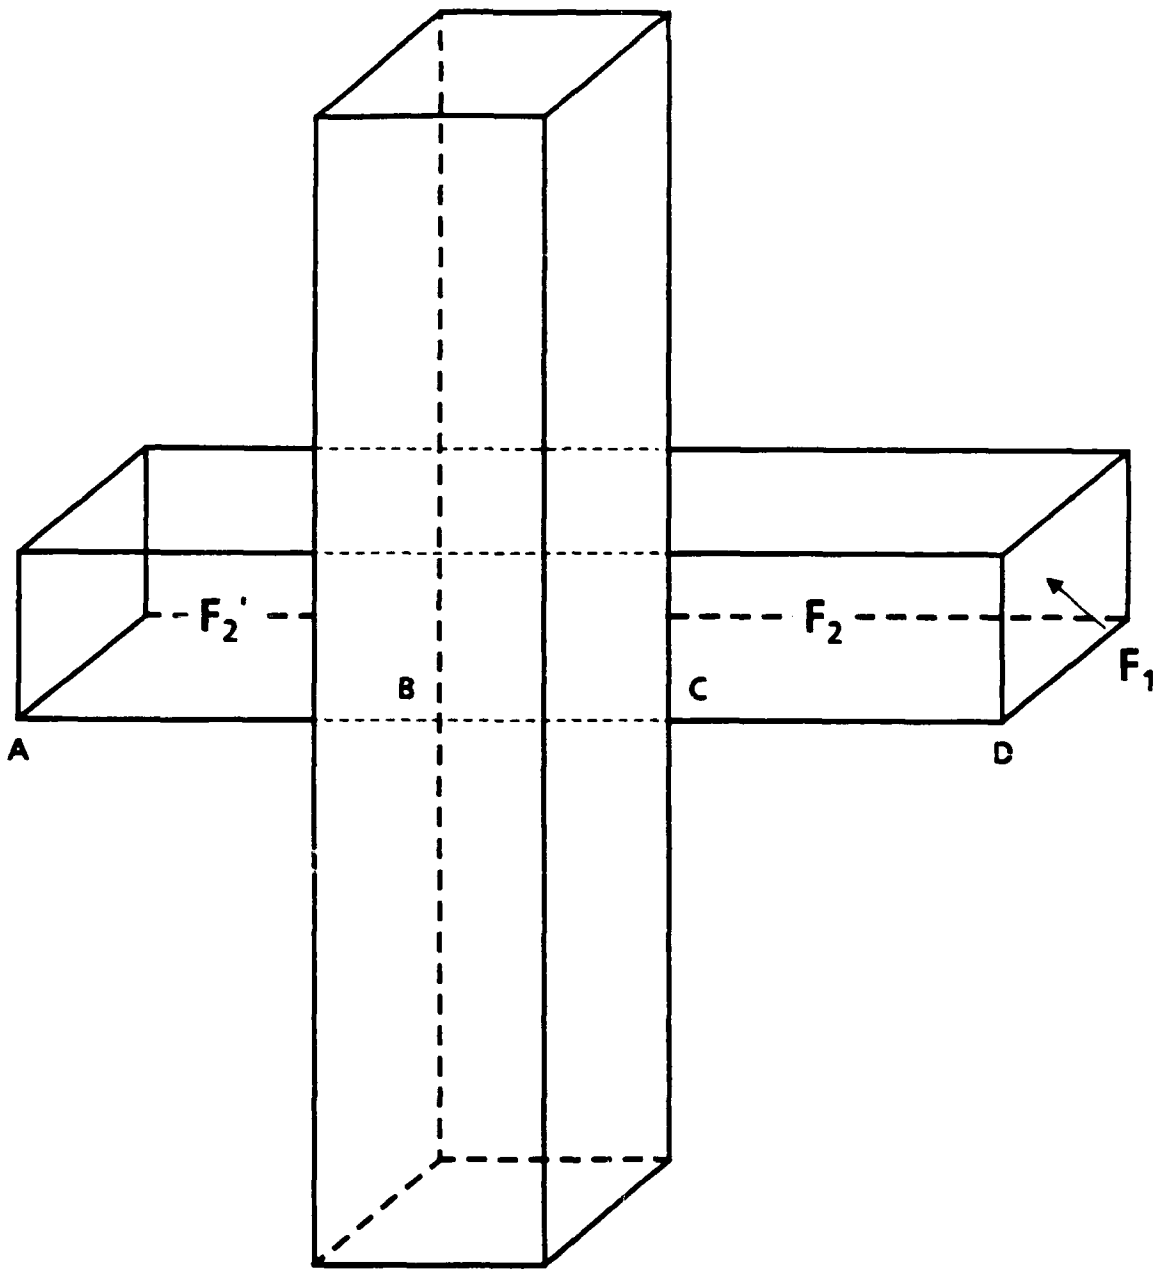

Figure 7: Two rectangular bars.

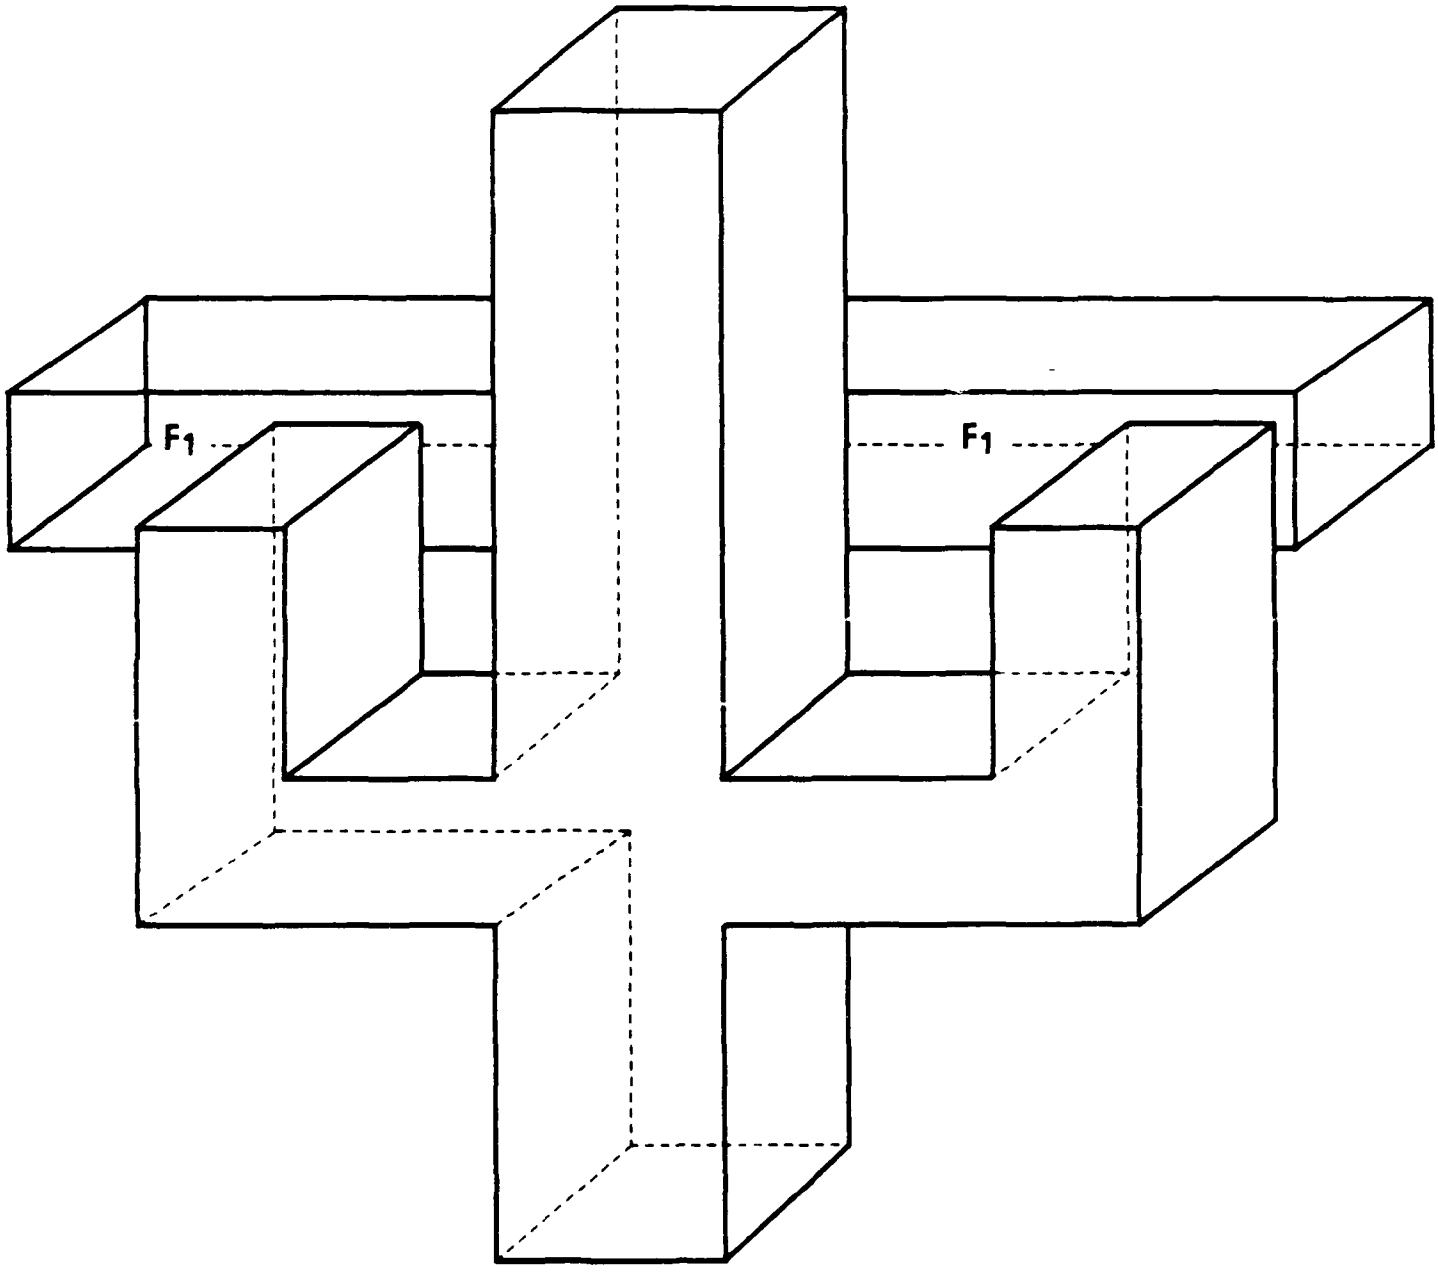

Figure 8: Force vertices on different faces to lie on the same line

Appendix A: Stressed graphs and polyhedra

This appendix provides proofs for Theorems 3.3.1, 3.3.4, and 3.6.1 (Corollary A.7, Proposition A.8, and Corollary A.10, respectively). We begin with a few elementary observations from linear algebra.

Every $v \in \mathbf{R}^3$ determines a linear function $v^*:\mathbf{R}^3 \to \mathbf{R}$ by the rule $v^*(w) = v \cdot w$, for every $w \in \mathbf{R}^3$. The triple product $[u,v,w] = \det(u,v,w)$ may then be expressed in terms of the cross product as $[u,v,w] = (v \times w)^*(u)$. The following lemma is immediate:

A.1 Lemma: $\{u_1, u_2, u_3\}$ is a basis of $\mathbf{R}^3$ if and only if

$$\{(u_2 \times u_3)^*, (u_3 \times u_1)^*, (u_1 \times u_2)^*\}$$

is a basis of $(\mathbf{R}^3)^*$; where $(\mathbf{R}^3)^*$ denotes the dual of $\mathbf{R}^3$, i.e., the vector space of all linear functions $\mathbf{R}^3 \to \mathbf{R}$. □

Let $\{u_1, u_2, u_3\}$ be a basis of $\mathbf{R}^3$. If $\lambda \in (\mathbf{R}^3)^*$, then there are unique constants c_1, c_2, and c_3 such that

$$\lambda = c_1(u_2 \times u_3)^* + c_2(u_3 \times u_1)^* + c_3(u_1 \times u_2)^*. \tag{1}$$

In fact, it is easy to see that c_i must be given by

$$c_i = \lambda(u_i)/\delta, \quad i = 1,2, \text{ and } 3 \tag{2}$$

where $\delta = [u_1, u_2, u_3]$.

Consider the vector space $A(H)$ of all affine functions $H \to \mathbf{R}$, where H is the hyperplane of $\mathbf{R}^3$ given by $z = 1$.

A.2 Lemma: $(\mathbf{R}^3)^*$ is isomorphic to $A(H)$ via the rule $\lambda \to \lambda|H$ (the restriction of λ to H).

Proof: The mapping $\lambda \to \lambda|H$ is clearly linear. Since $\lambda|H = 0$ implies that $\lambda = 0$ the maping is 1-1. It is onto because both $(\mathbf{R}^3)^*$ and $A(H)$ are 3-dimensional. □

Choose three noncollinear points p_1, p_2, and p_3 in H. As vectors in $\mathbf{R}^3$, they are linearly independent. Thus, Lemmas A.1 and A.2 yield the following two facts:

A.3 Corollary: For any $a \in A(H)$, there are unique constants c_1, c_2, and c_3 such that

$$a = c_1(p_2 \times p_3)^* + c_2(p_3 \times p_1)^* + c_3(p_1 \times p_2)^*|H. \tag{3}$$

In fact, setting $\delta = [p_1, p_2, p_3]$, it is easy to see that c_i must be given by

$$c_i = a(p_i)/\delta \quad , \quad i = 1, 2, \text{and } 3. \ \square \tag{4}$$

A.4 Corollary: Suppose a and a' are in $A(H)$, and $a(p_i) = a'(p_i)$ for $i = 1, 2$. Then there is a unique constant c such that

$$a - a' = c(p_1 \times p_2)^*|H. \ \square \tag{5}$$

It is instructive to evaluate c in this case. Choose any $p_* \in H$ not collinear with $\overline{p_1 p_2}$. Then, $(p_1 \times p_2)^*(p_*) = [p_1, p_2, p_*] \neq 0$. Therefore, after evaluating both sides of (5) at p_*, we may solve for c:

$$c = a(p_*) - a'(p_*)/[p_1, p_2, p_*]. \tag{6}$$

It follows that the expression on the right-hand side of (6) is independent of p_*.

Now recall that in §3 we defined $\Delta: PA(G) \to S(G)$ by setting $\Delta\{a_r\} = \{\omega_{ij}\}$, where, for each edge e_{ij}, ω_{ij} is given by

$$\omega_{ij} = \varepsilon(r,s)(a_s(p_*) - a_r(p_*))\,[p_i, p_j, p_*]. \tag{7}$$

Here f_r and f_s are the faces incident to e_{ij}, and $\varepsilon(r,s) = \pm 1$ according to our orientation convention, as described in §3. The preceding paragraph now implies the following result:

A.5 Corollary: ω_{ij} is independent of the choice of reference point p_r. □

A.6 Lemma: For any i, $\sum_{j=1}^{n} \omega_{ij}(p_i \times p_j) = 0$.

Proof: Let $f_1, f_2, \cdots, f_\ell$ denote the faces of $\Gamma(p)$ incident to p_i, in counterclockwise order around p_i. It will be convenient to write $r \rightarrow s$ if f_s is an immediate successor of f_r in this ordering. Choose any $r = 1, 2, \cdots, \ell$, and s such that $r \rightarrow s$. Then, choose $j \in \{1, 2, \cdots, n\}$ such that $f_r \cap f_s$ is the edge connecting p_i to p_j. Our orientation conventions are such that, if $r \rightarrow s$, then

$$\varepsilon(r,s) = \begin{cases} -1, & \text{for } j > i, \\ 1, & \text{for } j < i. \end{cases}$$

If $j > i$, then, combining (5), (6), and (7), we see that

$$\omega_{ij}(p_i \times p_j)^* \mid H = -(a_s - a_r) = a_r - a_s. \tag{8}$$

If $j < i$, then j and i in formula (7) get interchanged, and the right-hand denominator changes sign, as does $\varepsilon(r,s)$. Thus, we get equation (8) in this case too.

Now sum both sides of (8) as $r = 1, 2, \cdots, \ell$. The right-hand side telescopes to zero. Thus

$$\sum_{j=1}^{n} \omega_{ij}(p_i \times p_j)^* \mid H = 0, \tag{9}$$

from which the desired equation follows immediately via Lemma A.2. □

A.7 Corollary: $\{\omega_{ij}\}$ is an equilibrium stress on $\Gamma(p)$.

Proof: For any i and j, $p_i \times p_j = (p_i - p_j) \times p_i$. Therefore, by Lemma A.6,

$$\left(\sum_{j=1}^{n} \omega_{ij}(p_i - p_j)\right) \times p_i = \sum_{j=1}^{n} \omega_{ij}(p_i \times p_j) = 0,$$

from which it follows that $\sum_{j=1}^{n} \omega_{ij}(p_i - p_j)$ and p_i are linearly dependent. But, the former

belongs to $\mathbf{R}^2 = \{(x,y,0) \in \mathbf{R}^3\}$ whereas the latter belongs to $H = \{(x,y,1) \in \mathbf{R}^3\}$. So they can be linearly dependent if and only if

$$\sum_{i=1}^{n} \omega_{ij}(p_i - p_j) = 0 ,$$

as required. □

Theorem 3.3.4 asserts that, for every equilibrium stress $\{\omega_{ij}\}$, there is a P.A. sequence $\{a_r\}$ such that $\{\omega_{ij}\} = \Delta \{a_r\}$. We now give a proof of this fact. The proof makes use of some elementary concepts from homology theory. For a cell complex K and real vector space V, we shall assume that the reader is familiar with the concept of an n-chain on K with coefficients in V and the corresponding notions of boundary, cycle, and homology. We suggest [GH] or [HY] as a reference.

For our cell complex K, we use the 2-dimensional abstract complex given by $\Gamma(p)$. This consists of vertices p_i, edges e_{ij}, and faces f_r, with the incidence relations used throughout the paper. For our vector space V, we use $A(H)$, the vector space of all affine functions on H.

Typical 0-chains, 1-chains, and 2-chains will be denoted by

$$\sum a_i[p_i], \quad \sum b_{ij}[e_{ij}], \quad \text{and} \quad \sum c_r[f_r],$$

where we sum over all vertices, edges, and faces, respectively, and a_i, b_{ij}, and c_r are elements of $A(H)$.

Given any equilibrium stress $\{\omega_{ij}\}$ on $\Gamma(p)$, we now define a specific 1-chain $\sum b_{ij}[e_{ij}]$ by setting

$$b_{ij} = -\omega_{ij}(p_i \times p_j)^* \,|H.$$

Note that $b_{ij} = -b_{ji}$, for all i,j.

A.8 Lemma: $\sum b_{ij}[e_{ij}]$ is a 1-cycle.

Proof: For $\sum b_{ij}[e_{ij}]$ to be a 1-cycle its boundary must be zero. Formally, the boundary $\partial(\sum b_{ij}[e_{ij}])$ is a 0-chain $\sum a_i[p_i]$. Thus we must show that each $a_i = 0$. Assuming that each edge e_{ij} is oriented from p_i to p_j, the usual boundary formula yields, for $i = 1, 2, \ldots, n$,

$$a_i = -\sum_{j=i+1}^{n} b_{ij} + \sum_{j=1}^{i-1} b_{ji} = -\sum_{j=1}^{n} b_{ij} \tag{10}$$

$$= \sum_{j=1}^{n} \omega_{ij}(p_i \times p_j)^* \,|\, H.$$

Now we use the fact that $\{\omega_{ij}\}$ is an equilibrium stress:

$$\sum_{j=1}^{n} \omega_{ij}(p_i - p_j) = 0,$$

for each $i = 1, 2, \cdots, n$. Therefore, by reversing the argument of Corollary A.7, we get

$$\sum_{j=1}^{n} \omega_{ij}(p_i \times p_j) = 0,$$

which implies, using (10) above, that each $a_i = 0$, as desired. □

So far, all our results have been independent of the specific topology of K. We now make use of the fact that K is a cellular decomposition of the 2-sphere. It is well known that every 1-cycle on K is the boundary of a 2-chain. Thus, by Lemma A.9,

$$\sum b_{ij}[e_{ij}] = \partial(\sum a_r[f_r]), \tag{11}$$

for some 2-chain $\sum a_r[f_r]$, $a_r \varepsilon A(H)$. Here, ∂ is the usual boundary operator.

A.9 Proposition: For every equilibrium stress $\{\omega_{ij}\}$, there is a P.A. sequence $\{a_r\}$ such that $\{\omega_{ij}\} = \Delta\{a_r\}$.

Proof: Let an equilibrium stress $\{\omega_{ij}\}$ be given, and form the 1-chain $\sum b_{ij}[e_{ij}]$, as above. This is a 1-cycle, so as in (11), it is the boundary of a 2-chain $\sum a_r[f_r]$. We show that $\{a_r\}$ is a P.A. sequence and that $\Delta\{a_r\} = \{\omega_{ij}\}$.

Let e_{ij} be any edge. Formally, the coefficient of $[e_{ij}]$ in $\partial(\sum a_r[f_r])$ is $\varepsilon(r,s)(a_r - a_s)$, where f_r and f_s are the faces incident to e_{ij}.

Thus,

$$\varepsilon(r,s)(a_r - a_s) = b_{ij} = -\omega_{ij}(p_i \times p_j)^* \,|\, H. \tag{12}$$

Since the right-hand side vanishes at p_i and p_j, a_r and a_s coincide on e_{ij} Therefore, $\{a_r\}$ is a P.A. sequence.

Furthermore, if both sides of (12) are evaluated at p_* (which is noncollinear with p_i and p_j), we get

$$\omega_{ij} = \varepsilon(r,s)(a_s(p_*) - a_r(p_*)) \cdot [p_i, p_j, p_*],$$

which is precisely the formula used to define Δ. Therefore, $\Delta\{a_r\} = \{\omega_{ij}\}$, as desired. □

We conclude this appendix with a proof of the convexity criterion (Proposition 3.6.1). We begin with a surface S sitting on the plane H and corresponding P.A. sequence $\{a_r\}$, with $a_0 = 0$. Set $\{\omega_{ij}\} = \Delta\{a_r\}$. Clearly a necessary condition for the convexity of S is that the plane $z = 1$ be a supporting hyperplane. Thus, necessarily, S is contained either in the half-space $z \geq 1$ or in $z \leq 1$. By symmetry, we need deal only with the former case. Recall that, here, we always assume that the graph $\Gamma(p)$ is convex.

A.10 Lemma: Suppose S is contained in the half-space $z \geq 1$. Then the following are equivalent:

(a) S is strictly convex.

(b) For any pair of adjacent faces f_r and f_s, r and $s \geq 1$, and for any $q_r \in \mathit{interior}\ (f_r)$, we have

$$a_r(q_r) < a_s(q_r).$$

Proof: Let f_r and f_s be adjacent faces. Select q_r and q_s in the interiors of f_r and f_s respectively. Let $\ell(r,s) \subseteq H$ be the line segment connecting q_r to q_s: that is, $\ell(r,s) = \{tq_r + (1-t)q_s \mid 0 \leq t \leq 1\}$.

(a) ⇒ (b). Note that $a_r(q_r) \neq a_s(q_r)$, otherwise $a_r = a_s$, contradicting strict convexity. Thus, if (b) fails somewhere, we must have

$$a_r(q_r) > a_s(q_r), \tag{13}$$

for some r and s as in (b). Now note that (13) implies

$$a_r(q_s) < a_s(q_s). \tag{14}$$

To see this, restrict both a_r and a_s to $\ell(r,s)$, and consider the function values of $a_r(q)$ and

$a_s(q)$ as q ranges along $\ell(r,s)$ from q_r to q_s.

It follows from (13) and (14) that every value of

$$ta_r(q_r) + (1-t)a_s(q_s), \quad 0 < t < 1,$$

is strictly greater than either $ta_r(q_r) + (1-t)a_r(q_s) = a_r(tq_r + (1-t)q_s)$ or $ta_s(q_r) + (1-t)a_s(q_s) = a_s(tq_r + (1-t)q_s)$. Therefore, the line segment in $\mathbf{R}^3$ connecting $(q_r, a_r(q_r)) \varepsilon S$ to $(q_s, a_s(q_s)) \varepsilon S$ fails to lie in the region enclosed by S, contradicting convexity. Thus, (b) must hold.

(b) $\Rightarrow$ (a). The idea here is to show that if strict convexity is violated, there exist adjacent f_r and f_s etc., for which (b) fails. To find such f_r and f_s, we suppose that strict convexity is violated and find a line segment ℓ connecting points on S, interior to distinct faces, such that ℓ contains a third point on or above S. If all points of ℓ are on S, then S has coplanar adjacent faces which can serve as our f_r and f_s. Otherwise, choose $q \varepsilon \ell$ whose vertical distance above S is maximal and which projects vertically to an interior edge in G. The faces incident to this edge are the desired f_r and f_s. □

A.11 Corollary: S is strictly convex if and only if

$$\omega_{ij} < 0, \quad \text{for peripheral edges } e_{ij},$$

and

$$\omega_{ij} > 0, \quad \text{for interior edges } e_{ij}.$$

Proof: This follows immediately from Lemma A.9, together with formula (7) and the orientation conventions described in §3. □

Remark: A similar result is obtained if strict convexity is replaced by convexity and strict inequalities by weak inequalities.

Appendix B: A Generalization of Theorem 4.3.3

The theorem whose proof we sketch in this appendix generalizes Theorem 4.3.3 in that it applies to a more general convex polyhedron K. In Theorem 4.3.3, the convex polyhedron K is assumed to "sit over a triangular face" in the plane H, as described in §3. In the more general setting of this appendix, we allow K to sit over a face with any number of sides. In fact, even more generally, we consider any convex polyhedron K with no faces perpendicular to H, and we project it orthogonally to H. Although some faces of K may have overlapping images under this projection, each individual face is mapped in a 1-1 manner to H. As before, the 1-skeleton of K is isomorphic to an abstract graph Γ, and the projection of this 1-skeleton into H is a realization of Γ denoted $\Gamma(p)$. The periphery of $\Gamma(p)$ is a convex curve in H, and K sits over this graph as before. All of the results of §3 relating polyhedra to stressed graphs apply without change to this situation, which greatly generalizes that considered in Theorem 4.3.3. For example, the notion of a p.a. sequence $\{a_r\}$ is defined just as before, as is the notion of an equilibrium stress $\{\omega_{ij}\}$ for $\Gamma(p)$. And the correspondence $\{a_r\} <--> \{\omega_{ij}\}$ is defined just as in the more restricted setting in §3. With this understood, the theorem proved in this appendix now is a verbatim restatement of Theorem 4.3.3.

It is perhaps worth reminding the reader at this point why such a generalization is needed. First of all, we start with a plane P slicing a convex polyhedron K. If a triangular face of K has vertices a reasonable distance from P on each side of P, then there is no need for this generalization. However, this may not be the case. It may be that there are faces with vertices satisfying the condition just stated but none that is triangular. In that case, we normalize just as before. That is, we choose a central projection from a point on P near the said face (and outside of K), and project onto the plane of this face. We also apply a projective transformation so that this plane becomes H, P becomes the plane $x = 0$, and central projection becomes orthogonal projection onto H. The normalized K now sits over this non-triangular face, and the theorem of this appendix may be applied to it. But there remains a

further possibility. Namely, each face of K may be effectively on one side or the other of P. In this situation, we choose an edge both of whose endpoints are very close to P, and we perturb K slightly so that this edge is moved onto P. We then select one of the faces containing this edge and normalize as before. In this case, the projection point should be chosen on P and close enough to the face so that the resulting peripheral curve consists entirely of edges of the selected face and those of its neighbor. Again, the theorem of this appendix may now be applied to the normalized K.

We now start with K normalized as just described. Let $\bar{p}$ denote the n-tuple of vertices of the projection of K in H, and let $\bar{X}$ and $\bar{Y}$ denote the corresponding n-tuples of x- and y-coordinates, respectively. The stress constants $\bar{\omega}_{ij}$ corresponding to K determine a stress matrix $\bar{\Omega}$ as in §2, i.e., a symmetric $n \times n$ matrix with row-sums 0. Let e and f denote the number of edges and faces of K, respectively. Thus, $n - e + f = 2$. The spanning tree of $\Gamma(\bar{p})$ has $n - 1$ edges. Let us call the remaining $f - 1$ edges of $\Gamma(\bar{p})$ *free*. If we assign arbitrary stress constants to these edges, and we choose any n-tuple X suitably close to the n-tuple $\bar{X}$, then, as in Lemma 4.3.1, we may solve uniquely for the remaining stress constants so that x-equilibrium is satisfied. That is, we produce an $n \times n$ symmetric matrix Ω satisfying

$$\Omega \cdot X = 0. \tag{15}$$

Indeed, since Ω is a stress matrix, we also have

$$\Omega \cdot \bar{1} = 0, \tag{16}$$

where $\bar{1} = (1,1,\ldots,1)^t$ (cf. the proof of Lemma 2.3.1). For example, if $X = \bar{X}$ and the "free" stress constants are chosen to be the $\bar{\omega}_{ij}$ corresponding to free edges, then we obtain the matrix $\bar{\Omega}$. Let $\bar{s}$ denote the $(f-1)$-tuple of these free stress constants.

Our goal is to find an n-tuple Y, depending on X and an $(f-1)$-tuple of free stress constants s, such that, writing $Y = Y(X,s)$,

(a) $Y(X,s)$ is a continuous function of (X,s), with $Y(\bar{X},\bar{s}) = \bar{Y}$, and

(b) $$\Omega \cdot Y = 0. \tag{17}$$

For then, by restricting (X,s) close enough to $(\bar{X},\bar{s})$, we can force Y to be so close to $\bar{Y}$ that the corresponding graph $\Gamma(p)$ is convex. Furthermore, the matrix Ω can be forced to be close enough to $\bar{\Omega}$ to correspond to a convex stress on $\Gamma(p)$, which by (15)-(17) is an equilibrium stress. This produces the desired convex, stressed graph. The proof now concludes exactly as does that of Theorem 4.3.3.

To find a Y satisfying condition (1) above, we rely on the following result:

B.1 Proposition: rank $\bar{\Omega} = n-3$.

Proof : In the case of triangular periphery, this fact is an easy consequence of the results in §2. In any case, since $\bar{\Omega}$ satisfies (15)-(17) for $X = \bar{X}$, $Y = \bar{Y}$, it is clear that rank $\bar{\Omega} \leq n-3$. The proof of equality in general is similar to the proof of Theorem 5 of [C] and proceeds by contradiction. If rank $\bar{\Omega} < n$-3, then there exists a realization $\Gamma(q)$ of Γ in $\mathbf{R}^3$ projecting orthogonally onto $\Gamma(\bar{p})$, having 3-dimensional convex hull Q, and having $\{\bar{\omega}_{ij}\}$ as an equilibrium stress. In fact, Q is the convex hull of the vertices $q_1, q_2, \ldots, q_k$ that project onto the peripheral vertices $\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_k$ of $\Gamma(\bar{p})$. It is possible to find two such vertices that are connected by an edge e lying in the boundary of Q. The plane through e perpendicular to H divides Q into two halves. An infinitesimal rotation of one of these halves around e and towards the other half results in an ifinitesimal shortening of some of the "interior" edges of $\Gamma(q)$ without affecting the other "interior" edges. This shows that the energy function corresponding to Γ (cf. §2) does not have a critical point at q, contradicting Lemma 2.1.1. □

Now we relabel vertices if necessary so that the $n-3 \times n-3$ upper left corner $\bar{A}$ of $\bar{\Omega}$ is invertible. Then the same will hold for the upper left corner A of Ω, provided (X,s) is suitably close to $(\bar{X},\bar{s})$. Henceforth, let W_0 and W_1 denote suitably small neighborhoods of $\bar{X}$

and $\bar{s}$ in $\mathbf{R}^n$ and $\mathbf{R}^{r-1}$, respectively, and restrict (X,s) to $W_0 \times W_1$. Then we may write

$$\Omega = \begin{bmatrix} A & B \\ B^t & C \end{bmatrix}, \quad A \ (n-3)\times(n-3), \text{invertible.}$$

A simple computation shows that

$$\text{nullspace } \Omega \subseteq \text{columnspace} \begin{bmatrix} -A^{-1}B \\ I \end{bmatrix},$$

where I is the 3×3 identity matrix. Note that this column space is a 3-dimensional subspace of $\mathbf{R}^n$ which varies smoothly with (X,s) and coincides with nullspace $\bar{\Omega}$ when $(\bar{X},\bar{s}) = (X,s)$. In fact, the columnspace = nullspace Ω if and only if rank$\Omega = n-3$.

Define $Y = Y(X,s)$ to be the orthogonal projection of $\bar{Y}$ into the columnspace. Clearly this satisfies condition (1). Since condition (2) is equivalent to $Y \in$ nullspace Ω, the foregoing paragraph shows that we obtain condition (2) whenever rank$\Omega = n-3$. Thus, our proof will be complete as soon as we can insure this equality.

To formulate this more precisely, it is convenient to recognize that the spanning-tree procedure of Lemma 4.3.1 defines a smooth mapping

$$W_0 \times W_1 \xrightarrow{F} S(n),$$

where $S(n)$ is the vector space of all $n \times n$ stress matrices. The dimension of $S(n)$ is $\frac{1}{2}n(n-1)$. Let $S(n,r)$ denote the set of stress matrices of rank r. It is easy to show that $S(n,r)$ is a smooth submanifold of $S(n)$ of codimension $\frac{1}{2}(n-r)(n-r-1)$. We are interested in the codimension-three submanifold $S(n,n-3)$. By construction and Proposition B.1,

$$F(\bar{X},\bar{s}) = \bar{\Omega} \in S(n,n-3).$$

Let M denote the subset $F^{-1}(S(n,n-3)) \subseteq W_0 \times W_1$, and let $\pi: W_0 \times W_1 \to W_0$ be the pro-

jection map $\pi(X,s) = X$. The following is our main technical result.

B.2 Proposition: Assume that the graph $\Gamma(\bar{p})$ has a triangle (i.e., a simple, closed edge path consisting of three edges). Then

(a) There exists a neighborhood $V \subseteq W_0 \times W_1$ of $(\bar{X}, \bar{s})$, such that $M \cap V$ is a smooth manifold of dimension $n+f-2$.

(b) There exist open-ball neighborhoods B_0 around $\bar{X}$ in W_0 and B_1 around 0 in $\mathbf{R}^{f-2}$, together with a diffeomorphism $h: B_0 \times B_1 \to M \cap V$ satisfying

$$\pi h(x_1, \cdots, x_n, z_1, \cdots, z_{f-2}) = (x_1, \cdots, x_n),$$

for all $(x_1, \cdots, x_n, z_1, \cdots, z_{f-2})$ in $B_0 \times B_1$.

Let δ be the radius of B_0, and suppose $|\bar{X} - \lambda| < \delta$. Define s by the equation $(X,s) = h(X, 0)$. Then s depends smoothly on X, and $F(X,s)$ has rank n-3 by construction. Thus, assuming Proposition B.2 and the triangle hypothesis, we are done.

We now indicate the proof of Proposition B.2 and then how to get rid of the triangle hypothesis. Basically, it's a simple application of the Inverse (or Implicit) Function Theorem. The hard part is computing that a certain derivative is non-zero. We shall give a few details.

Rewrite Ω as

$$\begin{bmatrix} A & B_1 & B_2 \\ B_1^t & C_{11} & C_{12} \\ B_2^t & C_{12}^t & C_{22} \end{bmatrix} \begin{matrix} A & (n-3)\times(n-3), \text{invertible} \\ B_1 & (n-3)\times 2 \\ B_2 & (n-3)\times 1 \end{matrix}$$

B.3 Lemma: $\text{rank}\Omega = (n-3) + \text{rank}(C_{11} - B_1^t A^{-1} B_1)$. □

Thus, $(X,s) \in M$ provided $\text{rank}(C_{11} - B_1^t A^{-1} B_1) \neq 0$. Now $C_{11} - B_1^t A^{-1} B_1$ is a 2×2 symmetric matrix of rank ≤ 1 (because $\bar{1}$ and X belong to nullspace Ω). Such a matrix is zero $\Leftrightarrow$ it has zero trace. So define

$$g(X,s) = \text{trace}(C_{11} - B_1^t A^{-1} B_1).$$

This gives a real-valued function, $g: W_0 \times W_1 \to \mathbf{R}$, such that

$$F^{-1}(S(n, n-3)) = M = g^{-1}(0).$$

The advantage of g over F is that g is real-valued, and the kind of regularity result that we want is easier to obtain for g.

Now for any $X \in W_0$, define g_X to be g restricted to the "slice" $\{X\} \times W_1$: more precisely $g_X: W_1 \to \mathbf{R}$ is given by $g_X(s) = g(X, s)$. We then have the following result:

B.4 Proposition: Assume that the graph $\Gamma(\bar{p})$ has a triangle. Then $\bar{s}$ is a regular point of $g_{\bar{X}}$. That is, $d(g_{\bar{X}})_{\bar{s}} \neq 0$. Consequently, $(\bar{X}, \bar{s})$ is a regular point of $\pi \mid \mathbf{M} : \mathbf{M} \to W_0$.

Proposition B.2 follows easily from this.

The second assertion of Proposition B.4 follows from the first by standard arguments of differential topology. Therefore, it remains to prove $d(g_{\bar{X}})_{\bar{s}} \neq 0$. To show this, we need only compute this derivative in one judiciously chosen tangent direction. Such a tangent direction is an appropriate $(f-1)$-tuple of free stress constants. Now, with our triangle hypothesis, we are able to choose such constants so that the corresponding (infinitesimal) stress matrix has at most three non-zero rows and three non-zero columns. This greatly simplifies the computation and is the only reason for the hypothesis. The computation proceeds by cases according to the distribution of the selected rows and columns in the matrix. There are seven cases. We shall go through one case here for illustration.

We make a preliminary observation. A stress on Γ satisfying x-equilibrium with respect to $\bar{X}$ is uniquely determined by the stress assignments to the free edges. Thus, instead of selecting an assignment to the free edges, we select an entire stress. This is notationally convenient, since then we do not have to distinguish between stress constants that are free and those that are not free.

In the case that we consider here, the triangle of the hypothesis has no vertical edges and has vertices $\bar{p}_{n-2}, \bar{p}_{n-1}, \bar{p}_n$. We choose a stress satisfying x-equilibrium as follows. To all edges not in the triangle, assign stress constant 0. The edges $\bar{p}_{n-2}\bar{p}_{n-1}, \bar{p}_{n-1}\bar{p}_n, \bar{p}_{n-2}\bar{p}_n$ get the stress constants 1, a, b, respectively, where a and b are determined by the x-equilibrium requirement. It is not hard to check that $(1+a)(1+b) = 1$, again because of x-equilibrium. In this stress-direction, one calculates that the differential $d(g_{\bar{X}})_{\bar{S}}$ has the value

$$trace \begin{bmatrix} 1+b & -1 \\ -1 & 1+a \end{bmatrix} = a + b + 2,$$

which, by the preceding relation on a and b, has absolute value greater than two. Thus, Proposition B.4 is verified in this case. Complete details will appear in [K2].

This completes our discussion when $\Gamma(\bar{p})$ has a triangle.

In case $\Gamma(\bar{p})$ does not have a triangle, it must have an order-three vertex (in fact, at least eight of these). Perform a "slice" at this vertex, replacing it by a small triangle. Carry out the argument for the modified graph. Then restore the vertex by reversing the slice procedure.